

Lewis, Daniel J.; Mertens, Karel

Working Paper

Weak instrument bias in impulse response estimators

cemmap working paper, No. CWP01/26

Provided in Cooperation with:

The Institute for Fiscal Studies (IFS), London

Suggested Citation: Lewis, Daniel J.; Mertens, Karel (26) : Weak instrument bias in impulse response estimators, cemmap working paper, No. CWP01/26, Centre for Microdata Methods and Practice (cemmap), The Institute for Fiscal Studies (IFS), London, <https://doi.org/10.47004/wp.cem.2026.0126>

This Version is available at:

<https://hdl.handle.net/10419/339507>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Weak instrument bias in impulse response estimators

Daniel J. Lewis
Karel Mertens

The Institute for Fiscal Studies
Department of Economics, UCL

cemmap working paper CWP01/26



Weak Instrument Bias in Impulse Response Estimators*

DANIEL J. LEWIS

University College London

CEPR

KAREL MERTENS[†]

Federal Reserve Bank of Dallas

CEPR

January 2, 2026

Abstract

We approximate the finite-sample distribution of impulse response function (IRF) estimators that are just-identified with a weak instrument using the conventional local-to-zero asymptotic framework. Since the distribution lacks a mean, we assess bias using the mode and conclude that researchers prioritizing robustness against weak instrument bias should favor vector autoregressions (VARs) over local projections (LPs). Existing testing procedures are ill-suited for assessing weak instrument bias in IRF estimates, and we propose a novel simple test based on the usual first-stage F -statistic. We investigate instrument strength in several applications from the literature, and discuss to what extent structural parameters must be restricted ex-ante to reject meaningful bias due to weak identification.

JEL classification: C32, C36.

Keywords: Local Projections, Vector Autoregressions, Instrumental Variables, Weak Instruments, Impulse Responses, Dynamic Causal Effects.

*We thank seminar participants at the University of Aarhus, Board of Governors, Minneapolis Fed, University of Manchester, and Yale for helpful comments. We are grateful to Christiane Baumeister, Lutz Kilian, José Montiel Olea and Barbara Rossi for helpful discussions. All errors are ours. Funded in part by the European Union (ERC, CreMac, StG-101161596). Views and opinions expressed are those of the authors only and do not necessarily reflect those of the European Union, the European Research Council, the Federal Reserve Bank of Dallas or the Federal Reserve System. Neither the European Union nor the granting authority can be held responsible for the content of this paper.

[†]*Contact:* mertens.karel@gmail.com, tel: +(214) 922-6000, fax: +(214) 922-6302

Impulse response function (IRF) analysis plays a central role in empirical macroeconomics, particularly in estimating the dynamic effects of structural shocks. To facilitate economic interpretation of magnitudes and increase policy relevance, researchers frequently rescale impulse responses estimated with VARs or LPs to normalize the impact on specific endogenous variables. This scaling often means that, either explicitly or implicitly, impulse response estimates are functions of ratios of ordinary least squares (OLS) estimates.

As is well known in the literature on instrumental variables (IV), ratio-based estimators face particular statistical challenges in finite samples – see Andrews et al. (2019) for a recent overview. The estimators can be severely biased, and the ratio distribution can be far from the normal distribution dictated by central limit theorems. Macro applications may be particularly susceptible to such *weak instrument* problems, as samples are typically short and the fraction of plausibly exogenous variation in the scaling macro variables is often relatively small.

This paper develops new tools to analyze and address weak instrument bias in just-identified IRF estimators based on tractable approximations to their finite-sample distribution under the local-to-zero asymptotic framework of Staiger and Stock (1997). Since the resulting distribution lacks a mean, we measure bias using the major mode. This provides an intuitive and analytically convenient alternative to conventional mean-based definitions of bias, and it allows us to characterize how estimator behavior depends on the signal strength of the instrument, finite-sample endogeneity, and the number of horizons or endogenous variables considered.

Based on our analytical results, we derive three insights for applied work:

1. In the local-to-zero framework – which abstracts from all finite-sample distortions except for those induced by a weak instrument – the modal IRF bias in VARs never exceeds the bias in LPs, and is often smaller (for example, whenever the number of horizons equals or exceeds the number of variables in the VAR). Thus, when weak instrument bias is the primary concern, researchers should prefer VARs over LPs.
2. Testing procedures for weak instrument bias in the single-equation model (Stock and Yogo 2005; Montiel Olea and Pflueger 2013; Lewis and Mertens 2025) are not well-suited to most IRF applications. We describe

a simple test that is analogous to Stock and Yogo (2005) or Montiel Olea and Pflueger (2013), but uses the *modal* bias in *multiple* IRF estimates. The test utilizes the standard first-stage F -statistic, with critical values that are easily obtained with built-in commands in standard statistical software. We also describe extensions to allow for heteroskedasticity and serial correlation using the robust F -statistic.

3. Examining a range of IRF applications from the literature, our proposed test yields results that frequently diverge from existing procedures, and we conclude that weak identification often remains an important concern. We show that adding restrictions on structural parameters leads to a more powerful constrained weak instrument test, and we describe how to quantify the degree of restrictiveness required to rule out meaningful weak instrument bias in practical applications.

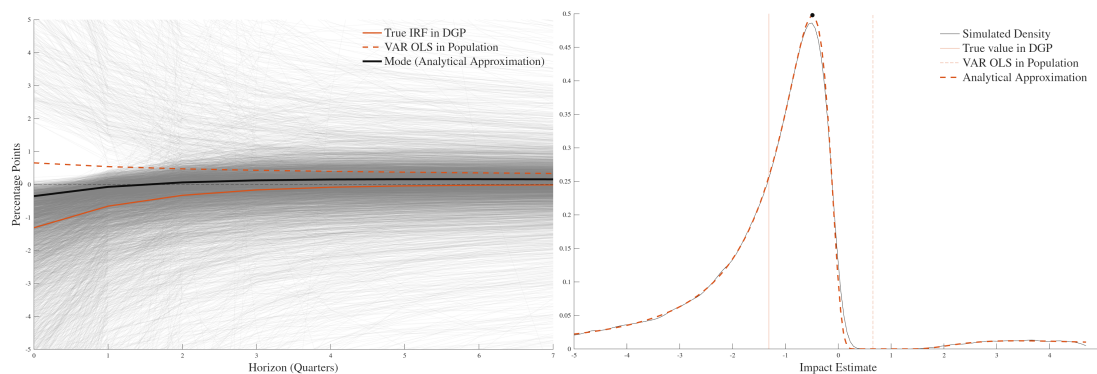
Our results are limited to the class of just-identified models (IRFs to a single shock identified using a single instrument), but are nevertheless relevant as the vast majority of applications fall into this category. We adopt a strictly frequentist perspective, though weak instruments pose very similar challenges for Bayesian approaches (Kleibergen and Zivot 2003; Giannone et al. 2025). While we focus on IRF estimation, our analytical results and testing procedures apply more broadly to just-identified system IV estimators that consider multiple outcome variables jointly but with a common first-stage. Finally, because the canonical just-identified single-equation model is nested within our framework, our procedure also provides a simple, novel weak instrument test with a sound theoretical foundation for that ubiquitous setting.

1 Weak Instrument Problems in IRF Estimation

1.1 An Illustrative Example

To illustrate the potential for weak instrument problems in macroeconomic applications, the left panel of Figure 1 plots estimates of the IRF of inflation to a one percentage point positive interest rate shock in samples of 250 quarters, generated from a simple New Keynesian model (see Appendix A for details). The estimated IRFs are from a VAR(1) model for the interest rate, inflation, and the output gap. To address the endogeneity of the interest rate,

Figure 1: SVAR-IV Impulse Responses of Inflation to an Interest Rate Shock



Notes: Estimated inflation responses to a one percentage point positive interest rate shock in Monte Carlo samples ($T = 250$) from a simple NK model with Gaussian shocks; see Appendix A. *Left:* SVAR-IV estimates with a valid instrument and a VAR(1) in the interest rate, output gap and inflation. “VAR OLS in Population” is the population IRF from a lower-triangular recursive ordering with the interest rate ordered first; “Mode (Analytical Approximation)” is the modal SVAR-IV response from Proposition 3. *Right:* “Simulated Density” is a kernel density estimate for the impact estimator based on 50,000 Monte Carlo samples. Analytical density and major mode (black circle) based on Proposition 3.

the responses are identified with random draws of an external instrument that – in population – has a correlation of 0.75 with the true exogenous interest rate shock in the model and is uncorrelated with the other structural shocks. By construction, the empirical VAR(1) model is fully consistent with the theoretical data-generating process (DGP) and the instrument is valid, so the only challenge for estimation is the small sample size of 250 quarters.

As the left panel in Figure 1 shows, the simulated estimates of the inflation responses exhibit substantial departures from a bell-shaped distribution centered on the model’s true IRF (solid orange line). These departures are especially pronounced in the tails: a significant fraction of the simulated IRF estimates show large persistent inflation increases following a positive interest rate shock, despite the model’s true response being negative at all horizons. Additionally, many IRF estimates exhibit much larger negative inflation responses than the true response in the model. The right panel of Figure 1 shows a kernel density estimate of the inflation responses in the first period only. As the figure shows, the density of the impact estimates is not centered on the true value (solid vertical line) but is heavily skewed, with the major mode –

the most probable estimate of the inflation impact – lying meaningfully to the right of the true impact in the theoretical model. The kernel density estimate further highlights the heavy tails and also shows a nontrivial range of values with no probability mass and a positive minor mode.

As is well known from the literature on weak instruments, the potential for nonstandard behavior as in Figure 1 is a general feature of ratio-based estimators (see, e.g., Nelson and Startz 1990). In the VAR model, the IRF estimates are linear combinations of the OLS coefficients from regressing the reduced form VAR residuals on the instrument, divided by the coefficient corresponding to a chosen scaling variable. The nonstandard behaviors result from two features of such ratio estimators in finite samples: (1) sampling uncertainty in the denominator and (2) nonzero correlation between the numerator and the denominator. Estimators of the denominator with positive density near zero induce heavy tails in the ratio distribution, while a nonzero correlation between the numerator and the denominator creates endogeneity problems by shifting the distribution of the ratio away from its population value.

The potential for weak instrument problems in IRF estimators is not specific to the use of a VAR or the external instrument-based identification approach, but exists for a broad range of widely-used IRF estimators. In IV-based LPs, for example, the IRF estimates are coefficients from regressing various leads of an endogenous variable on the instrument, scaled by the coefficient in the regression for the scaling variable.¹ In a VAR model with an internal instrument – which is a VAR expanded to include the instrument and Cholesky identification – the IRFs are also ratio-based whenever the responses are normalized by the impact on a variable other than the instrument. While the explicit use of instruments for IRF identification is now common in macroeconomic applications, many other existing approaches to IRF identification ultimately have an instrumental variables interpretation, and are therefore implicitly based on ratios of sample estimates as well (Hausman and Taylor 1983; Shapiro and Watson 1988).

Clearly, when judging empirical results based on any of these ratio-based IRF estimators in small samples, it is important to consider their potentially nonstandard properties.

¹The LP version of Figure 1 is provided in Appendix Figure A.1.

1.2 Addressing Weak Instrument Bias in IRFs

The primary focus of this paper is the impact of weak identification on the *bias* in IRF point estimates.² We also focus exclusively on the just-identified case, in which the dynamic responses to a single economic shock are identified with a single instrument. To obtain tractable analytical approximations to the finite-sample behavior of ratio-based IRF estimators, we follow the local-to-zero asymptotic framework of Staiger and Stock (1997). This framework assumes that the first-stage parameter – determining the instrument’s signal strength for the endogenous scaling variable – goes to zero at the rate $T^{-1/2}$ as the sample size T increases. As a consequence, both the numerator and the denominator in the ratio of the instrumental variable IRF estimator remain random variables even as T goes to infinity, as opposed to converging to their true population values as under conventional asymptotics with a fixed first-stage parameter. Assuming joint normality of the numerator and denominator based on standard central limit theorem arguments, we derive the density of the ratio. This density is well established when the numerator is a scalar (Fieller 1932; Marsaglia 1965; Hinkley 1969); it has been utilized to examine the finite-sample behavior of the *single*-equation IV estimator under just identification by Mariano and McDonald (1979) and Nelson and Startz (1990), among others. The IRF estimators studied in this paper are (linear transformations of) *multiple*-equation IV estimators, and therefore depend on ratios with vector-valued numerators. Accordingly, we extend the existing distribution theory to multivariate ratios of normal variables, in line with recent independent work by Yang and Gui (2023).

As in the just-identified univariate case, the joint distribution of the IRF coefficients is heavy-tailed and does not possess a finite mean or other positive-order moments. As a result, the conventional mean-based definition of estimator bias is inapplicable and an alternative measure of central tendency is required. We propose to assess bias using the (major) mode of the distribution. For heavy-tailed, asymmetric distributions, the mode provides a natural measure of noncentrality. It corresponds to the most likely realization of the estimator and thus serves as an intuitive and representative typical value for

²For IRF inference procedures that are robust to size distortions created by weak instruments, see Montiel Olea et al. (2021) and Jentsch and Lunsford (2022).

the point estimates encountered in empirical applications.³

Our primary analytical result is that, asymptotically, the difference between the modal IRF estimate and the true IRF is proportional to the bias of the OLS-based IRF estimate. The constant of proportionality is bounded between zero and one, and – for a given choice of IRF estimator – depends exclusively on three scalar parameters implied by the underlying DGP: (1) the *concentration parameter* (the standardized noncentrality of the denominator); (2) the *endogeneity parameter* (the squared canonical correlation between the numerator and the denominator); and (3) the *rank of the IRF* (defined as the minimum of the ratio estimator’s vector dimension and the number of response horizons). The concentration parameter determines the probability mass near the true population value, and can be interpreted as a measure of the effective sample size. The endogeneity parameter determines the shift in probability mass towards the population OLS estimates. The rank of the IRF determines the dimensionality of the ratio distribution.

The proportionality result enables direct, simulation-free comparisons of modal bias across estimators and relative to OLS or other benchmarks that depend only on these three parameters. We leverage this result to answer three questions for applied work:

Question 1: Local Projections or Vector Autoregressions?

Plagborg-Møller and Wolf (2021) show that, under suitable specifications, LPs and VARs estimate the exact same impulse responses in population and differ only in their finite-sample behavior. Recent studies of the finite-sample behavior of LP- and VAR-based IRF estimators have focused mainly on the bias-variance trade-off in finite samples, and specifically, the tension between lag truncation bias in VARs and the higher estimation variance inherent in LPs (Schorfheide 2005; Li et al. 2024; Montiel Olea et al. 2025).⁴ Our paper adds a distinct perspective by considering the finite-sample bias arising from a weak instrument.

³An alternative is to measure bias using the median, as in Angrist and Kolesár (2024) and Lewis and Mertens (2025). However, the median lacks the analytical tractability of the mode and, in multivariate settings, does not have a universally accepted definition.

⁴Others focus on finite-sample inference in LPs vs. VARs, (e.g. Kilian and Kim 2011; Montiel Olea et al. 2024).

Under the general conditions of Plagborg-Møller and Wolf (2021), LPs and VARs share a common estimand, which implies that – in the local-to-zero framework – they also have the same asymptotic OLS bias. Provided the first-stage specifications are identical, LP and VAR-based IV IRF estimators also share the same concentration parameter. Our proportionality result implies that any difference in the modal IV bias must therefore stem solely from the endogeneity parameter and the rank of the IRF. We prove analytically that comparing these two parameters yields an unambiguous ranking: the VAR bias never exceeds the LP bias, regardless of the norm used to aggregate the biases across individual IRF coefficients. Moreover, apart from special cases, the VAR bias is generally smaller than the bias in the corresponding LP. In practice, this implies that applied researchers primarily concerned with weak instrument bias should favor VARs over LPs.

We emphasize that the weak dominance of VARs in terms of modal IV bias holds strictly within the local-to-zero asymptotic framework. While this framework provides analytical tractability by isolating weak identification issues, it abstracts from other finite-sample distortions, such as bias arising from VAR lag misspecification. As Plagborg-Møller and Wolf (2021) demonstrate, if the VAR lag length is misspecified and smaller than in the true DGP, VARs and LPs share the same estimand only for horizons up to the lag length. Consequently, our result – that VAR asymptotic bias does not exceed LP asymptotic bias – is guaranteed only for these horizons. Beyond the lag length, researchers face a trade-off between reduced weak instrument bias and the risk of lag truncation bias.⁵

Question 2: How to Test for Weak Instrument Bias in IRFs?

To assess the potential for weak instrument bias in IV-identified IRFs, researchers commonly report results of tests for weak instruments, typically based on a first-stage F -statistic. In line with suggestions in Montiel Olea et al. (2021), researchers usually rely on the conventional rule-of-thumb threshold of 10 for the critical value (Staiger and Stock 1997; Stock and Yogo 2005), or alternatively on the critical values in Montiel Olea and Pflueger (2013).

⁵High persistence in time series data is another potentially significant source of finite-sample bias that the local-to-zero asymptotic framework abstracts from. We are not aware of any general ranking of LP and VAR bias in this context (e.g., Herbst and Johannsen 2024).

These testing procedures are ill-suited for most IRF applications for two reasons. First, existing weak instrument tests are for specifications with a single outcome variable, while IRF analyses typically aim to estimate dynamic responses over multiple horizons. When using LPs, one can test each horizon separately. But in many cases, a more appropriate approach is to treat the IRF as a single causal estimand and combine the bias across all horizons of interest to conduct a joint test. Moreover, for VARs, the existing tests are only applicable for the impact horizon, but cannot be applied to detect bias at other horizons.

Second, the vast majority of instrumental variable IRF applications are just identified, with responses to a single shock estimated using a single instrument. In such settings, the ratio estimator lacks a finite mean. Because standard weak instrument tests are predicated on mean bias, they are theoretically invalid for just-identified models, as also made clear by Stock and Yogo (2005).

To address these shortcomings of existing tests in the context of IRF estimation, we describe a simple novel test based on the usual first-stage F -statistic.⁶ The test’s null hypothesis is analogous to that of Montiel Olea and Pflueger (2013): the IV bias relative to the bias in the most-adverse scenario of perfect endogeneity and an irrelevant instrument exceeds a user-specified tolerance level, τ . Without any restrictions on the underlying DGP, the test’s null hypothesis is also analogous to that of Stock and Yogo (2005): the IV bias relative to the OLS bias exceeds tolerance τ . The only difference relative to both existing tests is that the bias criterion considers the modal IV bias rather than the mean IV bias. Because of the proportionality of the mode to the asymptotic bias of OLS, the critical value for the test is remarkably straightforward to obtain and – unlike the mean-based test for over-identified models in Stock and Yogo (2005) – does not require numerical simulations. Table 1 in Section 4.2 provides critical values for conventional choices of the tolerance level, τ , the significance level, α , and a range of values for the rank of

⁶In unpublished work, Lunsford (2016) proposes a weak instrument test for just-identified instrumental variable VARs based on the bias in the contemporaneous impact coefficients. As in Stock and Yogo (2005), critical values controlling the mean bias are obtained from Monte Carlo simulations. Because the distribution lacks a mean, we regard the simulated critical values as unreliable. In addition, the test is restricted to impact responses and does not extend to LPs. Angelini et al. (2024) propose a weak instrument test for VAR IRFs identified indirectly by instruments for all non-targeted shocks. While appealing, their test is not applicable to the vast majority of applications of IV-based IRF estimators.

the IRF.⁷ We also provide a bootstrap refinement for obtaining critical values, a constrained test that incorporates ex-ante restrictions on structural parameters, and extensions to accommodate heteroskedasticity and serial correlation.

Question 3: Is Weak IV Bias a Concern in IRF Applications?

We evaluate concerns about weak instrument bias in four VAR studies employing instrumental variable IRF estimators spanning a range of different applications (monetary policy, oil supply, uncertainty, and tax policy shocks). While all four studies yield F -statistics exceeding the conventional rule-of-thumb of 10 – and most exceed the more stringent threshold in Montiel Olea and Pflueger (2013) – our test frequently fails to reject the null hypothesis of weak instruments, so weak instrument bias often remains a concern.

As is the case in the existing single-equation tests, our baseline test is designed to protect against the bias under the worst-case scenario for the true values of the unknown structural parameters. In our setting, this worst-case DGP always features arbitrarily large dynamic causal effects and a vanishingly small variance contribution of the targeted shock. Our recommendation to practitioners is to look for plausible restrictions on the structural parameters to restrict the set of admissible DGPs to exclude such extreme scenarios.

Across the four applications, we evaluate the minimum required variance contribution of the targeted shock to assess the degree of restrictiveness required to reject the null of a weak instrument. In many cases, this null can be rejected either without restrictions, or else with reasonably small minimum variance requirements. One exception is the identification of monetary policy shocks via a high-frequency instrument, following Gertler and Karadi (2015). Within their specific empirical framework, rejection requires assuming that endogeneity bias is small from the outset, indicating that the high-frequency instrument provides limited incremental value for the reliable identification of monetary policy shocks.

⁷For the practical determination of the rank of the IRF in VARs and LPs, see Section 2.3.

2 Estimating IRFs with an Instrument: Analytical Results

2.1 Modeling Framework

The object of interest is the vector of impulse responses $\delta \in \mathbb{R}^H$. To nest different IRF estimators (VARs and LPs), we express δ as a general linear combination of another parameter vector, $\beta \in \mathbb{R}^N$:

$$(1) \quad \delta = \theta + \Theta\beta,$$

where $\theta \in \mathbb{R}^H$ and $\Theta \in \mathbb{R}^{H \times N}$ and $\text{rank}(\Theta) = \min\{H, N\}$. The parameters in β arise from a system of N simultaneous structural equations:

$$(2) \quad \forall t: \quad y_t = \beta Y_t + u_t, \quad \mathbb{E}[u_t] = 0,$$

where $y_t \in \mathbb{R}^N$ is a vector of N outcome variables and $Y_t \in \mathbb{R}$ is correlated with the structural errors u_t . An instrumental variable $z_t \in \mathbb{R}$ satisfies the exogeneity condition $\forall t: \mathbb{E}[z_t u_t] = 0$.

When $H = 1$ and $\delta = \beta$, the above framework reduces to the standard just-identified single-equation IV model. The general modeling framework is an $N \times 1$ system of equations identified using the same single instrumental variable. While the framework's application is not restricted to impulse response estimation, several widely used IRF estimators fit into this setting:

- **LP-IV**: Consider (lag-augmented) local projections of x_t^y at horizons $h = 0, \dots, H - 1$ on x_t^Y with predetermined controls X_{t-1} :

$$(3) \quad x_{t+h}^y = \beta_h x_t^Y + \gamma_h X_{t-1} + u_{h,t},$$

where z_t instruments for x_t^Y to identify β_h . Here, Y_t is the residual from projecting x_t^Y on X_{t-1} ; y_t is the $H \times 1$ vector of residuals from projecting x_{t+h}^y on X_{t-1} ; and u_t stacks the $u_{h,t}$ terms. The vector β contains the H impulse response coefficients, so $\delta = \beta$ and $N = H$, $\theta = 0$ and $\Theta = \mathcal{I}_H$.

- **SVAR-IV (external)**: Let x_t^y contain $K - 1$ variables and consider

$x_t = [x_t^Y, x_t^{y'}]' \in \mathbb{R}^K$ with the structural VAR(p) representation

$$(4) \quad x_t = \sum_{j=1, \dots, p} A_j x_{t-j} + e_t, \quad e_t = D \epsilon_t, \quad D = \begin{bmatrix} d_Y & d'_{Yy} \\ \beta d_Y & D_y \end{bmatrix},$$

where $e_t = [e_t^Y, e_t^{y'}]'$ are one-step-ahead forecast errors, ϵ_t are the structural shocks, and $D \in \mathbb{R}^{K \times K}$ is the impact matrix. Define $\Gamma = [\mathcal{I}_{Kp} : \mathcal{A}' : (\mathcal{A}^2)' : \dots : (\mathcal{A}^{H-1})']$ where \mathcal{A} is the companion matrix, and let \mathbf{e}_n be the n -th column of \mathcal{I}_{Kp} . Defining

$$(5) \quad \begin{aligned} \theta' &= [1 : \mathbf{0}_{1 \times Kp-1}] \Gamma (\mathcal{I}_H \otimes \mathbf{e}_n), \\ \Theta' &= [\mathbf{0}_{K-1 \times 1} : \mathcal{I}_{K-1} : \mathbf{0}_{K-1 \times K(p-1)}] \Gamma (\mathcal{I}_H \otimes \mathbf{e}_n), \end{aligned}$$

$\delta = \theta + \Theta \beta$ contains the first H impulse responses of the n -th variable to the first structural shock, normalized to a unit effect on x_t^Y . In applications, z_t is typically an external measure or proxy for the unobserved structural shock (external instrument or Proxy SVAR). In this setup, $Y_t = e_t^Y$, $y_t = e_t^{y'}$, $N = K - 1$ and β is the ratio of the contemporaneous impact of the identified shock on x_t^y to that on x_t^Y .

- **SVAR-IV (internal)**: An internal-instrument SVAR augments the VAR in (4) to $\tilde{x}_t = [\tilde{z}_t, x_t^Y, x_t^{y'}]' \in \mathbb{R}^{\tilde{K}}$ with $\tilde{K} = K + 1$ and applies a recursive identification scheme. In this approach, Y_t and y_t are the residuals in the projections of x_t^Y and x_t^y on $\tilde{x}_{t-1}, \dots, \tilde{x}_{t-p}$, and the instrument z_t is the residual in the projection of \tilde{z}_t on $\tilde{x}_{t-1}, \dots, \tilde{x}_{t-p}$. The definitions of β , θ , and Θ above can be easily extended to the augmented VAR, and $N = \tilde{K} - 1 = K$. Alternative recursive orderings are possible; as long as the normalizing variable x_t^Y differs from the internal instrument \tilde{z}_t , the model above applies.

The general formulation in (1)-(2) highlights the fundamental distinction between LPs and VARs. In the LP framework, the H impulse response coefficients are treated as direct parameters within a system of equations for the leads of the outcome variable. Conversely, in VARs, these coefficients are derived from the underlying autoregressive parameters, A_1, \dots, A_p , and the $K - 1$ contemporaneous impact parameters, β , which govern the dynamics of K endogenous variables. Consequently, because LP-IV and SVAR-IV estimators

rely on distinct structural systems, their finite-sample properties diverge.

Unless stated otherwise, we continue to interpret δ as containing the first H impulse responses of a single variable for clarity. However, stacking responses of multiple endogenous variables, selecting arbitrary horizon subsets, or analyzing derived objects—such as cumulative, weighted, or discounted impulse responses—can be handled by suitably redefining x_t^y , θ , and Θ .

To ensure analytical tractability, we focus exclusively on the finite-sample challenges in estimating β and rely otherwise on asymptotic approximations. Accordingly, we treat the auxiliary parameters – the γ_h 's in LPs and A_j 's in VARs – as fixed in the analytical derivations. This simplification is justified asymptotically provided that consistent estimators are available for these auxiliary parameters. We state this high-level assumption explicitly as follows:

Assumption 1. *In LPs, a consistent estimator $\hat{\gamma}_h \xrightarrow{p} \gamma_h$ is available for $h = 0, \dots, H - 1$. In VARs, a consistent estimator $\hat{A}_j \xrightarrow{p} A_j$ is available for $j = 1, \dots, p$.*

In VARs, this assumption and the continuous mapping theorem guarantee that θ and Θ are also consistently estimated. Unless mentioned otherwise, the constructed terms y_t, Y_t , and z_t , as well as the parameters θ and Θ , are henceforth treated as known.

2.2 The Instrumental Variable IRF Estimator

Consider the first-stage equation

$$(6) \quad \forall t: \quad Y_t = \Pi z_t + w_t, \quad \text{E}[w_t] = 0, \quad \text{E}[z_t w_t] = 0,$$

where Π is the scalar first-stage parameter. Without loss of generality, we let $\text{E}[z_t] = 0$ and $\text{E}[z_t^2] = 1$. Substituting (6) in (2) yields the reduced form

$$(7) \quad y_t = \beta \Pi z_t + v_t,$$

where $v_t = \beta w_t + u_t$ such that $\text{E}[v_t] = 0$ for all t .

Let $y \in \mathbb{R}^{T \times N}$ and $Y, z \in \mathbb{R}^T$ collect a sample of T observations of the variables y_t, Y_t and z_t for $t = 1, \dots, T$. Let $Z = z(z'z/T)^{-\frac{1}{2}}$ such that $Z'Z/T = 1$. With this normalization, the OLS estimators of the scalar first-stage parame-

ter Π and the $N \times 1$ reduced form parameters $\Pi\beta$ are, respectively

$$(8) \quad \widehat{\Pi} = Y'Z/T, \quad \widehat{\Pi}\widehat{\beta} = y'Z/T.$$

The IV estimator of the impulse response vector δ is

$$(9) \quad \widehat{\delta}_{IV} = \theta + \Theta\widehat{\beta}_{IV}, \quad \text{where } \widehat{\beta}_{IV} = y'Z/(Y'Z) = \widehat{\Pi}\widehat{\beta}/\widehat{\Pi}.$$

The latter is simply the vector of OLS estimates of the reduced-form parameters in (7) divided by the scalar OLS estimate of the first-stage parameter in (6). Below, we will regularly compare the IV estimator, $\widehat{\delta}_{IV}$, to the OLS estimator, $\widehat{\delta}_{OLS} = \theta + \Theta\widehat{\beta}_{OLS}$, which is the IRF implied by the OLS estimate $\widehat{\beta}_{OLS} = y'Y/(Y'Y)$. In a VAR, $\widehat{\delta}_{OLS}$ also coincides with the IRFs to a unit innovation in the normalizing variable obtained with a lower triangular recursive identification scheme and the normalizing variable ordered first (Montiel Olea et al. 2021; Plagborg-Møller and Wolf 2021).

Let $w = Y - Z\Pi$ and $u = y - Y\beta'$. Since $w'Z/T \xrightarrow{p} E[z_t w_t] = 0$, under standard (strong-instrument) asymptotics it follows that $\widehat{\Pi}$ is consistent for Π , i.e. $\widehat{\Pi} \xrightarrow{p} \Pi$. Under exogeneity of the instrument, $u'Z/T \xrightarrow{p} E[z_t u_t] = 0$, such that $\widehat{\Pi}\widehat{\beta} \xrightarrow{p} \Pi\beta$. Assuming the instrument is relevant ($\Pi \neq 0$), then $\widehat{\beta}_{IV} \xrightarrow{p} \beta$, and consequently $\widehat{\delta}_{IV} \xrightarrow{p} \delta$. In other words, as T increases, $\widehat{\delta}_{IV}$ converges to the true value δ , providing the asymptotic justification for using the IV-based IRF estimator to identify causal effects.

2.3 The IRF Estimator with a Weak Instrument

As the illustrative example in Figure 1 shows, the behavior of the IRF estimator, $\widehat{\delta}_{IV}$, can be highly nonstandard in finite samples because of weak identification. Following Staiger and Stock (1997), we consider an asymptotic approximation of the finite-sample distribution of $\widehat{\delta}_{IV}$ with a weak instrument. Specifically, we replace the first-stage equation in (6) with

$$(10) \quad Y_t = \Pi_T z_t + w_t,$$

where

Assumption 2. $\Pi_T = c/\sqrt{T}$, where c is a fixed scalar constant.

Under this assumption, the first-stage coefficient is local to zero. This ensures that sampling uncertainty remains significant relative to Π_T even in the limit ($T \rightarrow \infty$), effectively embedding finite-sample concerns within an asymptotic structure. Assumption 2 implies that $\widehat{\Pi} \xrightarrow{p} 0$ and $\widehat{\Pi}\beta \xrightarrow{p} 0$, so that the IV estimator $\hat{\beta}_{IV}$ – the ratio of $\widehat{\Pi}\beta$ to $\widehat{\Pi}$ – does not converge to β as $T \rightarrow \infty$. Consequently, the IRF estimator, $\hat{\delta}_{IV}$, is now inconsistent.

To characterize the distribution of $\hat{\delta}_{IV}$ under local-to-zero asymptotics, we impose two additional high-level assumptions. First,

Assumption 3.

$$T^{-1} \begin{bmatrix} w & u \end{bmatrix}' \begin{bmatrix} w & u \end{bmatrix} \xrightarrow{p} \Sigma_{wu} = \begin{bmatrix} \sigma_w^2 & \sigma'_{wu} \\ \sigma_{wu} & \Sigma_u \end{bmatrix}, \quad \Sigma_{wu} \text{ positive definite.}$$

This assumption ensures that the asymptotic covariance matrix of the error terms is full-rank and requires the system in (2) to be stochastically non-singular, a conventional assumption in both VAR and LP models. One minor implication is that, if the outcome variable of interest is the normalizing variable itself, the impact horizon must be excluded from δ because it introduces an identity in (2).⁸

The second additional high-level assumption is that a central limit theorem applies for the first-stage and reduced form regression scores,

Assumption 4. *Let $w = Y - Z\Pi$ and $v = w\beta' + u$; then*

$$T^{-\frac{1}{2}} \begin{bmatrix} w'Z \\ v'Z \end{bmatrix} \xrightarrow{d} \begin{bmatrix} \eta_1 \\ \eta_2 \end{bmatrix} \sim \mathcal{N}(\mathbf{0}_{(N+1) \times 1}, \Sigma_{wv}),$$

where $T^{-1} \begin{bmatrix} w & v \end{bmatrix}' \begin{bmatrix} w & v \end{bmatrix} \xrightarrow{p} \Sigma_{wv} = \begin{bmatrix} \sigma_w^2 & \sigma'_{wv} \\ \sigma_{wv} & \Sigma_v \end{bmatrix}$, $\Sigma_v = \Sigma_u + \beta\beta'\sigma_w^2 + \sigma_{wu}\beta' + \beta\sigma'_{wu}$ and $\sigma_{wv} = \sigma_w^2\beta + \sigma_{wu}$.

Assumption 3 guarantees that the covariance matrix Σ_{wv} is positive definite. Besides asymptotic normality, Assumption 4 also imposes conditional homoskedasticity and serially uncorrelated regression scores (CHSU). The CHSU assumption implies that Σ_{wv} coincides with the asymptotic covariance matrix

⁸Assumption 3 also prevents setting β equal to δ in VARs by redefining y_t as $\theta Y_t + \Theta y_t$, since Σ_{wu} would be singular whenever $H > N = K - 1$.

of the OLS estimates $\widehat{\Pi}$ and $\widehat{\Pi}\beta$. In Section 4.3.3 below, we relax the CHSU assumption and discuss its implications.

Defining $\nu_1 = c + \eta_1$ and $\nu_2 = \eta_2 - \beta\eta_1$, Assumptions 2-4 imply that

$$(11) \quad \widehat{\beta}_{IV} - \beta \xrightarrow{d} \frac{\nu_2}{\nu_1}, \quad \begin{bmatrix} \nu_1 \\ \nu_2 \end{bmatrix} \sim \mathcal{N} \left(\begin{bmatrix} c \\ \mathbf{0}_{N \times 1} \end{bmatrix}, \Sigma_{wu} \right)$$

In words, the bias $\widehat{\beta}_{IV} - \beta$ converges to a random vector that is a quotient of normally distributed variables, where the numerator ν_2 is an $N \times 1$ vector and the denominator ν_1 is a scalar.

Two scalar parameters are key for the shape of the ratio distribution:

$$(12) \quad \mu^2 = c^2/\sigma_w^2 \geq 0, \quad 0 \leq \rho^2 = \sigma'_{wu}\Sigma_u^{-1}\sigma_{wu}/\sigma_w^2 < 1.$$

The first parameter, μ^2 , is the *concentration parameter*, which governs the extent to which probability mass is concentrated near the origin. The second parameter, ρ^2 , is the *endogeneity parameter*, which is the squared canonical correlation between w and u , also equal to the asymptotic R^2 from regressing w on u . If the true model is an SVAR as in (4), $1 - \rho^2$ also measures the contribution of the identified shock to the variance of Y . The endogeneity parameter governs the extent to which the IV estimator is pulled toward the asymptotic OLS bias, with larger values implying greater endogeneity bias. Positive definiteness of Σ_{wu} (Assumption 3) ensures that $0 \leq \rho^2 < 1$.

We first consider the special case of an irrelevant instrument, $\mu^2 = 0$. In this case, the denominator of $\widehat{\beta}_{IV}$ is centered at zero, and the asymptotic distribution is a multivariate t -distribution with one degree of freedom or, equivalently, the multivariate Cauchy distribution.

Lemma 1. *When $\mu^2 = 0$,*

$$\widehat{\beta}_{IV} - \beta \xrightarrow{d} \frac{\nu_2}{\nu_1} \sim \mathcal{C} \left(\frac{\sigma_{wu}}{\sigma_w^2}, \frac{\Sigma_{u|w}}{\sigma_w^2} \right),$$

where

$$\Sigma_{u|w} = \Sigma_u - \frac{\sigma_{wu}\sigma'_{wu}}{\sigma_w^2} = \Sigma_v - \frac{\sigma_{wv}\sigma'_{wv}}{\sigma_w^2},$$

and $\mathcal{C}(m, \Sigma)$ denotes the multivariate Cauchy distribution with density

$$\varphi_{\mathcal{C}}(b, m, \Sigma) = \frac{\Gamma\left(\frac{N+1}{2}\right)}{\pi^{(N+1)/2}(\det(\Sigma))^{1/2}} \left[1 + (b - m)' \Sigma^{-1} (b - m)\right]^{-\frac{N+1}{2}},$$

where $\Gamma(s) = \int_0^\infty q^{s-1} e^{-q} dq$ is the gamma function.

The location vector, σ_{wu}/σ_w^2 , and the scale matrix, $\Sigma_{u|w}/\sigma_w^2$, match the mean and covariance of the asymptotic distribution of $\hat{\beta}_{OLS}$ under the local-to-zero assumption.⁹ Thus, with an irrelevant instrument, the asymptotic distribution of $\hat{\beta}_{IV}$ is centered at the asymptotic OLS value, $\beta + \sigma_{wu}/\sigma_w^2$, where σ_{wu}/σ_w^2 is the OLS bias. Owing to the heavy tails of the Cauchy distribution, the mean and all positive-order moments are undefined.

The next result extends Lemma 1 to the general case, $\mu^2 \geq 0$.

Proposition 1. *Under Assumptions 1-4, $\hat{\beta}_{IV} - \beta \xrightarrow{d} \nu_2/\nu_1$ with density*

$$\varphi(b) = \varphi_{\mathcal{C}}(b, \sigma_{wu}/\sigma_w^2, \Sigma_{u|w}/\sigma_w^2) \cdot h(\zeta(b)),$$

where

$$\zeta(b) = \frac{|\mu|}{1 - \rho^2} (\sigma'_{wu} \Sigma_u^{-1} b - 1) \left(1 + (b - \sigma_{wu}/\sigma_w^2)' (\Sigma_{u|w}/\sigma_w^2)^{-1} (b - \sigma_{wu}/\sigma_w^2)\right)^{-\frac{1}{2}};$$

$$h(t) = \frac{e^{\frac{1}{2} \left(t^2 - \frac{\mu^2}{1 - \rho^2}\right)}}{\Gamma\left(\frac{N+1}{2}\right)} \cdot \begin{cases} \sum_{j=0,2,4,\dots}^N \binom{N}{j} (t/\sqrt{2})^{N-j} \Gamma\left(\frac{j+1}{2}\right) & \text{if } N \text{ even;} \\ -\sum_{j=0}^N \binom{N}{j} (-|t|/\sqrt{2})^{N-j} \gamma\left(\frac{j+1}{2}, \frac{t^2}{2}\right) \\ \quad + \sum_{j=1,3,5,\dots}^N \binom{N}{j} (t/\sqrt{2})^{N-j} \Gamma\left(\frac{j+1}{2}\right) & \text{if } N \text{ odd;} \end{cases}$$

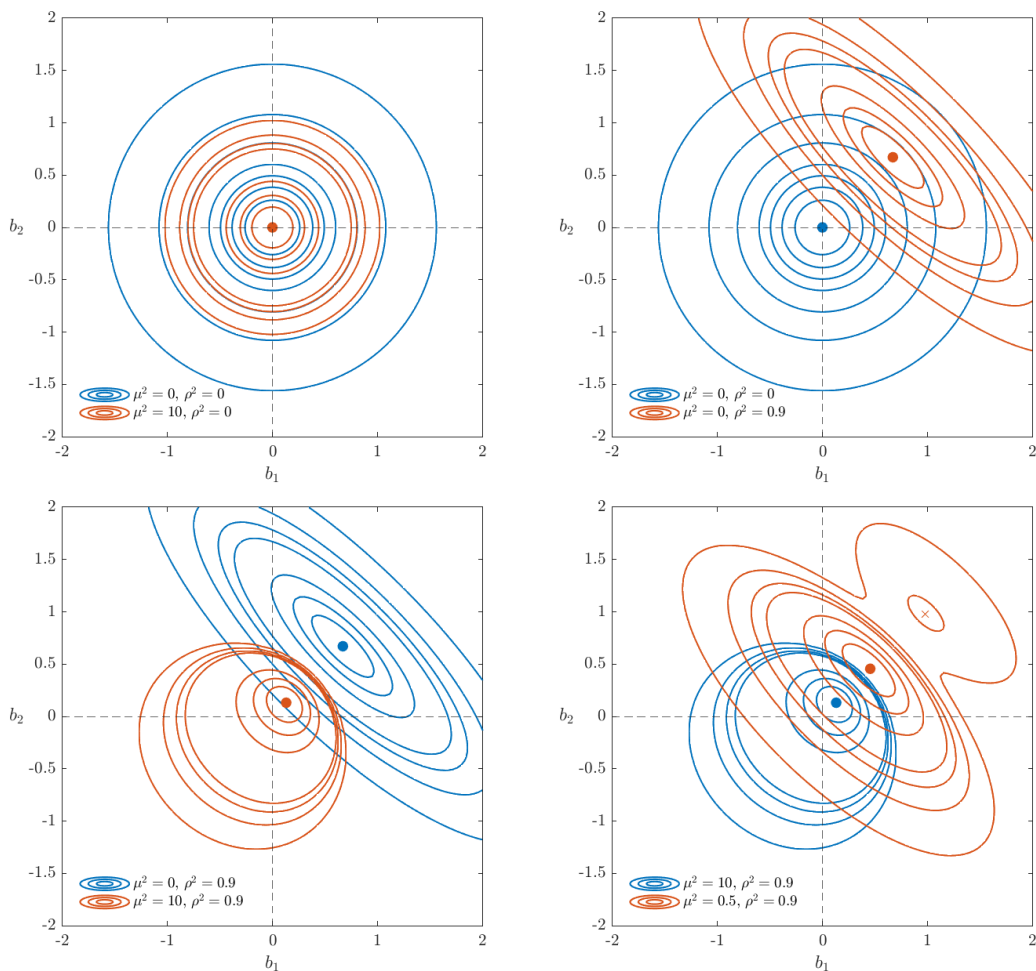
$\Gamma(s) = \int_0^\infty q^{s-1} e^{-q} dq$ is the gamma function and $\gamma(s, v) = \int_0^v q^{s-1} e^{-q} dq$ is the lower incomplete gamma function.

Proof. See Appendix B.1. □

Proposition 1 nests Lemma 1 as a special case when $\mu^2 = 0$.¹⁰ When $\mu^2 > 0$, the first-stage contains some signal, which changes the shape of the distribution through the multiplicative factor $h(\zeta(b))$ in the density. The concentration parameter, μ^2 , governs the extent to which probability mass is concentrated on $b = 0$, while the endogeneity parameter ρ^2 determines how far the distribution

⁹Under Assumptions 1-3 and a central limit theorem, $\sqrt{T}(\hat{\beta}_{OLS} - \beta) \xrightarrow{d} \mathcal{N}(\sigma_{wu}/\sigma_w^2, \Sigma_{u|w}/\sigma_w^2)$.
¹⁰When $\mu^2 = 0$, $\zeta(b) = 0$ and $h(\zeta(b)) = 1$ for all $b \in \mathbb{R}^N$.

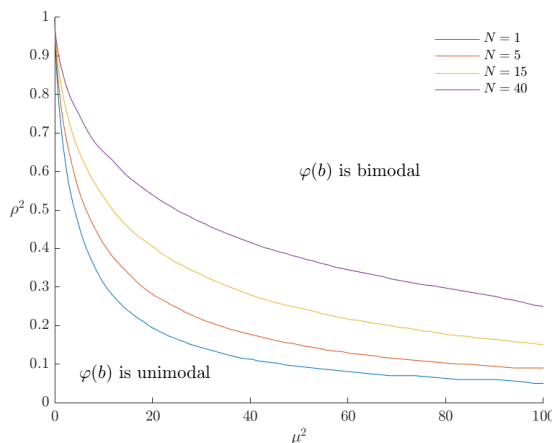
Figure 2: Examples of $\varphi(b)$ for $N = 2$



Notes: Contours of $\varphi(b)$ for various values of μ^2 and ρ^2 . All panels are for $\sigma_w^2 = 1, \Sigma_u = \mathcal{I}_2$, and a 45 degree angle of σ_{wu} . Circles (crosses) denote major (minor) modes. The distribution for $\mu^2 = 10$ and $\rho^2 = 0.9$ in the lower left and right panels has a minor mode to the northeast that lies outside of the ranges plotted.

is shifted toward the OLS bias, as illustrated for $N = 2$ in Figure 2. At $b = 0$, the exponential term in $h(\zeta(b))$ remains constant while the polynomial terms grow with μ^2 , concentrating more probability mass at the truth. For $b \neq 0$, the exponential term declines with μ^2 , reducing probability mass away from $b = 0$. Higher values of ρ^2 , on the other hand, shift mass in a direction determined by σ_{wu} , that is, towards the OLS bias. Importantly, even for $\mu^2 > 0$, the distribution remains heavy-tailed: since $\lim_{\|b\| \rightarrow \infty} h(\zeta(b)) = h(|\mu\rho|/\sqrt{1-\rho^2}) > 0$,

Figure 3: Modality Regions of $\varphi(b)$.



Notes: $\varphi(b)$ has one mode below the lines shown, and two above.

the tails decay at the same rate as a Cauchy distribution, implying that the mean and all positive-order moments are undefined.

When $\mu^2 = 0$, the unique mode of $\varphi(b)$ coincides with the asymptotic OLS bias σ_{wu}/σ_w^2 (Lemma 1). For $\mu^2 > 0$, the distribution is not necessarily unimodal, but can exhibit bimodality depending on the values of μ^2 and ρ^2 , as illustrated in the lower right panel of Figure 2. Figure 3 displays the modality regions as a function of μ^2 and ρ^2 for various dimensions N . Figures 2 and 3 illustrate that the qualitative properties of $\varphi(b)$ for arbitrary N remain similar to the $N = 1$ case discussed by Nelson and Startz (1990).

The following proposition characterizes the major mode of $\varphi(b)$:

Proposition 2. *The major mode of $\varphi(b)$ in Proposition 1 is*

$$b_{IV}^+ = \arg \max \varphi(b) = \frac{\sigma_{wu}/\sigma_w^2}{1 + s_N^+},$$

where $s_N^+ = (1 - \rho^2)(k^+ - 1) \geq 0$, $k^+ = \lim_{n \rightarrow \infty} k_n$, and $\{k_n\}$ is generated by the recurrence

$$k_{n+1} = \frac{1}{2} \left(g(r(k_n)) + \sqrt{g(r(k_n))^2 + \frac{4\rho^2}{1 - \rho^2} (g(r(k_n)) - 1)} \right),$$

$$r(k) = -\frac{k|\mu|(1 - \rho^2)^{-1/2}}{\sqrt{k^2 + \rho^2/(1 - \rho^2)}}, \quad g(t) = 1 + \frac{1}{N + 1} \frac{dh(t)}{dt} \frac{t}{h(t)},$$

with initial condition $r(k_{-1}) = |\mu|(1 - \rho^2)^{-1/2}$. For $\mu^2 > 0$, $1/(1 + s_N^+)$ is strictly decreasing in ρ^2 and μ^2 and strictly increasing in $N = \dim(\beta)$.

Proof. See Appendix B.2. □

Proposition 2 shows that, for $T \rightarrow \infty$, the most probable value of $\hat{\beta}_{IV} - \beta$ is proportional to the asymptotic OLS bias, σ_{wu}/σ_w^2 . The proportionality constant, $1/(1 + s_N^+)$, which lies between zero and one, decreases with the concentration parameter μ^2 and the endogeneity parameter ρ^2 and increases with the dimension $N = \dim(\beta)$, but is otherwise independent of other model parameters. Although there is no closed-form expression for the proportionality constant, Proposition 2 provides a simple and efficient iterative algorithm that converges to the global maximum of $\varphi(b)$ with arbitrary accuracy.

The next proposition establishes how the asymptotic distribution of $\hat{\beta}_{IV}$ translates to that of the impulse response estimator $\hat{\delta}_{IV}$.

Proposition 3. *Define the rank of the IRF, $R = \text{rank}(\Theta) = \min\{H, N\}$. Under Assumptions 1-4, $\hat{\delta}_{IV} - \delta \xrightarrow{d} \Theta\nu_2/\nu_1$ with density*

$$\varphi_\delta(d) = \begin{cases} \tilde{\varphi}(d), & \text{if } R = H \leq N; \\ (\det(\Theta'\Theta))^{-1/2} \varphi(\Theta^\dagger d), & \text{if } R = N < H; \end{cases}$$

where $\varphi_\delta(d)$ is defined with respect to the Hausdorff measure \mathcal{H}^N for \mathcal{H}^N -almost every $d \in \{\Theta b : b \in \mathbb{R}^N\}$ when $R = N < H$, $\Theta^\dagger = (\Theta'\Theta)^{-1}\Theta'$ is the Moore-Penrose inverse, and $\tilde{\varphi}(\cdot)$ is the pdf in Proposition 1 after replacing N , Σ_{wu} , and ρ^2 with, respectively,

$$R, \quad \tilde{\Sigma}_{wu} = \begin{bmatrix} \sigma_w^2 & \sigma'_{wu}\Theta' \\ \Theta\sigma_{wu} & \Theta\Sigma_u\Theta' \end{bmatrix}, \quad \tilde{\rho}^2 = \sigma'_{wu}\Theta'(\Theta\Sigma_u\Theta')^{-1}\Theta\sigma_{wu}/\sigma_w^2.$$

The major mode of $\varphi_\delta(d)$ is

$$d_{IV}^+ = \operatorname{argmax} \varphi_\delta(d) = \begin{cases} d_{OLS}/(1 + \tilde{s}_R^+) & \text{if } R = H < N; \\ \Theta b_{IV}^+ = d_{OLS}/(1 + s_R^+) & \text{if } R = N; \end{cases}$$

where $d_{OLS} = \Theta\sigma_{wu}/\sigma_w^2$, s_R^+ is defined as in Proposition 2 with N replaced by R , and \tilde{s}_R^+ is defined as in Proposition 2 with N and ρ^2 replaced by R and $\tilde{\rho}^2$,

respectively.

Proof. See Appendix B.3. □

In the LP case, $\hat{\delta}_{IV} = \hat{\beta}_{IV}$, so Propositions 1 and 2 apply directly, providing both the distribution and the major mode of $\hat{\delta}_{IV}$ under local-to-zero asymptotics. When $\hat{\delta}_{IV} = \theta + \Theta \hat{\beta}_{IV}$, as in VAR applications, the distribution additionally depends on Θ as described in Proposition 3. Importantly, the proportionality result continues to hold for $\hat{\delta}_{IV}$: the modal IV bias is proportional to d_{OLS} , which is the IRF bias implied by the population value of the OLS estimate of β , given by $\beta + \sigma_{wu}/\sigma_w^2$. As an illustration, the analytical mode and distribution for the New Keynesian example of Section 1.1 are depicted in Figure 1.

The mode of the asymptotic distribution of $\hat{\delta}_{IV}$ depends on R , the rank of IRF. The following examples illustrate the practical determination of R in LP and VAR settings.

- **LP.** For impulse responses estimated via LP-IV, the rank of the impulse response equals the number of elements in $\hat{\delta}_{IV}$. For the response of a single variable over H horizons, $R = H$. For impulse responses of multiple outcome variables, R equals the number of variables multiplied by H .
- **VAR.** In internal or external SVAR-IV, the rank of the impulse response is given by $R = \min\{H, K - 1\}$, where K is the total number of endogenous variables in the VAR, regardless of the lag length of the VAR. For example, in a four-variable VAR with impulse responses over six horizons, $R = \min\{6, 4 - 1\} = 3$. In a ten-variable VAR with four horizons, $R = \min\{4, 10 - 1\} = 4$. The rank, R , remains unchanged even when stacking the responses of multiple endogenous variables.

Thus, imposing VAR restrictions on the dynamics reduces the rank of IRF, R , from H to $\min\{H, K - 1\}$, where K denotes the number of endogenous variables in the VAR. Unlike in LPs, R does not scale with the total number of impulse response coefficients but is bounded by the number of endogenous variables in the VAR system.

The proportionality of the mode to the asymptotic OLS bias in Proposition 3 is central to the results in the remainder of this paper, as it facilitates analytical comparisons of modal IV bias across LP and VAR estimators, as

well as assessments of bias relative to the benchmarks in the criteria of first-stage tests like that of Stock and Yogo (2005) or Montiel Olea and Pflueger (2013).

3 Comparing weak instrument Bias in LPs and VARs

An important practical question for researchers is whether to use VAR or LP models for IRF estimation. To help guide that choice, the following proposition establishes a general ranking of the modal IV bias in just-identified VAR- and LP-based IRF estimators in the local-to-zero asymptotic framework:

Proposition 4. *Let $\hat{\delta}_{IV}^{LP}$ and $\hat{\delta}_{IV}^{VAR}$ be just-identified instrumental variable IRF estimators of $\delta \in \mathbb{R}^H$ in LP and VAR models, respectively. Suppose Assumptions 1-4 hold, and that the LP and VAR models share the same first-stage equation and OLS impulse response estimand. Then, for any norm $\|\cdot\|$,*

- (i) $\|d_{IV}^{+LP}\| \geq \|d_{IV}^{+VAR}\|$;
- (ii) $\|d_{IV}^{+LP}\| > \|d_{IV}^{+VAR}\|$ if $H \geq K$, $\rho_{VAR}^2 > 0$, and $\mu_{VAR}^2 > 0$;
- (iii) $\|d_{IV}^{+LP}\| > \|d_{IV}^{+VAR}\|$ if $\rho_{VAR}^2 > \rho_{LP}^2$, and $\mu_{VAR}^2 > 0$;

where K is the number of endogenous variables in the VAR, and d_{IV}^{+LP} and d_{IV}^{+VAR} are defined as in Proposition 3.

Proof. See Appendix B.4. □

Proposition 4 establishes that – under the stated assumptions – the modal IV bias in the VAR asymptotically never exceeds the modal IV bias in LPs for all conventional definitions of the magnitude of a vector. Apart from trivial cases (an irrelevant instrument or no OLS endogeneity bias), the VAR bias is strictly smaller whenever the number of impulse response coefficients H equals or exceeds the number of endogenous variables K included in the VAR. The VAR bias is also strictly smaller whenever its endogeneity parameter, ρ_{VAR}^2 , is strictly larger than the endogeneity parameter in the LP system, ρ_{LP}^2 – regardless of the number of impulse response coefficients.

Plagborg-Møller and Wolf (2021) show that, under suitable specifications, LP and VAR models target identical impulse response estimands under quite

general conditions.¹¹ Given Assumptions 1-4, the OLS-based IRF estimators, $\hat{\delta}_{OLS}^{LP}$ and $\hat{\delta}_{OLS}^{VAR}$, are consistent; consequently, both converge to the common estimand and display the same OLS bias, d_{OLS} , as $T \rightarrow \infty$ under those conditions. The IV-based IRF estimators, $\hat{\delta}_{IV}^{LP}$ and $\hat{\delta}_{IV}^{VAR}$, in contrast, are not consistent in the local-to-zero framework. They converge in distribution to random vectors with different densities, as they are based on distinct systems of structural equations. Including the same controls in the LP and VAR first-stage equations ensures that the concentration parameters are identical in both densities, $\mu_{LP}^2 = \mu_{VAR}^2 = \mu^2$. When d_{OLS} and μ^2 are identical for both densities, Proposition 3 implies that comparisons of the modal IV bias can be based solely on comparisons of the rank of the IRF, R , and the endogeneity parameter, ρ^2 . The rankings in Proposition 4 follow from the fact that $R_{LP} \geq R_{VAR}$, with strict inequality in (ii), and $\rho_{VAR}^2 \geq \rho_{LP}^2$, with strict inequality in (iii), and the fact that the constant of proportionality in the mode is strictly decreasing in ρ^2 and strictly increasing in R for $\mu^2 > 0$ (see Proposition 2).

The determination of R in VAR and LP models, and the fact that $R_{LP} \geq R_{VAR}$, was already discussed in the previous section. The proof that $\rho_{VAR}^2 \geq \rho_{LP}^2$ is in Appendix B.4. Intuitively, ρ^2 measures the explanatory power of the structural errors u for the first-stage residual w . By assumption, the first-stage residual is identical in LPs and the VAR, and also asymptotically equal to Y in the local-to-zero framework. The explanatory power of the structural errors for Y in a VAR model (contemporaneous errors across all of the VAR variables other than Y) is always at least as large than that of the structural errors in an LP system (cross-horizon errors for only the subset of variables of interest). Consequently, the endogeneity parameter is weakly larger in the VAR. That the VAR has smaller modal IV bias because its endogeneity parameter is larger sounds counterintuitive. Note, however, that endogeneity affects both the proportionality constant and the OLS bias in the expression for the mode. The latter is by assumption identical in the comparison, and the smaller VAR bias arises because the proportionality constant decreases with ρ^2 .

Proposition 4 provides a broad justification to prefer VARs over LPs in

¹¹In recent work, Ludwig (2024) establishes a finite sample equivalence between LPs and VARs by expressing LPs as combinations of equal and higher order VARs and VARs as combinations of equal and lower order LPs. In our setting, this results in equal structural systems and identical IRF estimators.

applications where weak instrument bias is the dominant concern for IRF estimation. It is important to note, however, that the advantage of VARs in terms of modal IV bias is derived strictly within the local-to-zero asymptotic framework. This framework focuses on the weak instrument problem, but ignores all other sources of finite-sample bias. Specifically, Plagborg-Møller and Wolf (2021) show that when the VAR lag length is insufficient, VAR and LP estimands align only up to that specific horizon. In that case, our finding that VARs exhibit weakly lower modal IV bias is only asymptotically guaranteed for horizons within the lag length of the VAR. At longer horizons, a trade-off arises between mitigating weak instrument bias and potentially incurring truncation bias.

4 Testing for Weak Instrument Bias in IRF Estimators

Another important practical question for researchers is how to assess the potential for weak instrument bias in IRF applications. Researchers currently mostly default to existing tests for single-equation models based on a first-stage F -statistic (Stock and Yogo 2005; Montiel Olea and Pflueger 2013; Lewis and Mertens 2025). However, these tests are only suitable for assessing weak instrument bias in individual IRF coefficients. Moreover, in VARs they are only applicable to assess bias in the impact coefficients, which are typically not the (only) main objects of interest. A second problem is that single-equation tests usually target the mean of the local-to-zero asymptotic distribution. Most applications of instrumental-variable IRF estimators in the literature employ just-identified models, where the mean of the distribution does not exist.¹² Next, we propose a test targeting the modal IV bias across multiple estimates jointly, making it better suited for assessing weak instrument bias in IRF applications.

¹²Stock and Yogo (2005) only tabulate critical values for models with degrees of overidentification of two or more. Researchers often resort to a rule-of-thumb critical value of 10 for the first-stage F -statistic, or to the critical values in Montiel Olea and Pflueger (2013) controlling the mean bias in a second-order (Nagar) approximation, a metric that in just-identified models lacks a clear relationship to the actual (infinite) bias it aims to approximate. The test in Lewis and Mertens (2025) targets the median in just-identified models, but the median has no universally accepted definition in multivariate settings.

4.1 Bias Criterion for the Test

As in existing single-equation tests, the goal is to construct a statistical test of the null hypothesis that the instrument is weak, with a precise definition of “weak” in terms of the potential IRF bias under the local-to-zero asymptotic framework. As discussed above, we consider the bias in terms of the (major) mode, d_{IV}^+ . For brevity, we henceforth assume $H > R = N$. The formulas that follow also apply to the $R = H \leq N$ case after substituting s_R^+ and ρ^2 with \tilde{s}_R^+ and $\tilde{\rho}^2$ as in Proposition 3.

The bias criterion for the test is as follows:

Definition 1. *The bias criterion for the impulse response estimator $\hat{\delta}_{IV}$ is*

$$B = \sqrt{\frac{d_{IV}^{\prime\prime} W^\dagger d_{IV}^+}{1 + d_{OLS}^{\prime\prime} W^\dagger d_{OLS}}} = \frac{|\rho|}{1 + s_R^+},$$

where W^\dagger is the Moore-Penrose inverse of the asymptotic covariance of the OLS-implied impulse response, $\hat{\delta}_{OLS} = \theta + \Theta \hat{\beta}_{OLS}$, given by

$$W = \begin{cases} \tilde{\Sigma}_{u|w}/\sigma_w^2; & \text{if } R = H \leq N; \\ \Theta \Sigma_{u|w} \Theta' / \sigma_w^2; & \text{if } H > R = N; \end{cases}$$

and d_{IV}^+ and s_R^+ are defined as in Proposition 3.

The bias criterion, B , is determined by a quadratic loss in the modal IV bias, $d_{IV}^{\prime\prime} W^\dagger d_{IV}^+ = \rho^2(1 - \rho^2)^{-1}(1 + s_R^+)^{-2}$. The weighting by the Moore-Penrose inverse of the asymptotic OLS covariance, W^\dagger , ensures the criterion remains invariant to arbitrary re-scalings of the outcome variable, y , or the endogenous regressor, Y . To constrain the criterion between zero and one, we scale the quadratic IV loss term. Specifically, we normalize by $1 + d_{OLS}^{\prime\prime} W^\dagger d_{OLS} = (1 - \rho^2)^{-1}$ rather than by the quadratic OLS loss term itself, $d_{OLS}^{\prime\prime} W^\dagger d_{OLS} = \rho^2(1 - \rho^2)^{-1}$. This results in $B = |\rho|(1 + s_R^+)^{-1}$, and since $0 \leq |\rho| < 1$ and $s_R^+ \geq 0$, it follows that $0 \leq B(\rho^2, \mu^2) < 1$. The choice of the scaling factor is driven by the behavior of the ratio with respect to ρ^2 . As shown in Proposition 2, the ratio $d_{IV}^{\prime\prime} W^\dagger d_{IV}^+ / (d_{OLS}^{\prime\prime} W^\dagger d_{OLS})$ is equal to $(1 + s_R^+)^{-1}$, which is decreasing in ρ^2 . Consequently, scaling directly by the OLS loss would yield a bias criterion that decreases as ρ^2 increases and is maximized

at $\rho^2 = 0$, even though both the IV and OLS bias terms individually are strictly increasing in ρ^2 . By scaling by $1 + d'_{OLS}W^\dagger d_{OLS}$ instead, we ensure that B is increasing in $\rho^2 \in [0, 1)$ for all but a small range of concentration parameter values. Specifically, B increases in ρ^2 for $\mu^2 \in [\underline{\mu}^2, \infty) \cup \{0\}$, where $\underline{\mu}^2 = 2 \left(\sqrt{1 + (R+1)^2} - (R+1) \right) > 0$.¹³ The bias criterion, B , depends only on three parameters: the endogeneity parameter ρ^2 , the concentration parameter μ^2 , and the rank of the IRF, R . Since R is known, we will use the notation $B(\rho^2, \mu^2)$ whenever it is instructive to highlight the dependence on the unknown parameters.

The units of the bias criterion can be interpreted in several ways. The bias criterion approaches its upper bound of unity when the instrument is irrelevant ($\mu^2 = 0$) and as Y becomes perfectly correlated with u ($\rho^2 \rightarrow 1$). Therefore, as in Montiel Olea and Pflueger (2013), one possible interpretation of $B(\rho^2, \mu^2)$ is the bias relative to the most adverse scenario of an irrelevant instrument and an endogenous regressor that is perfectly correlated with the structural errors. For $\rho^2 \rightarrow 1$, the criterion also shares the same interpretation as in Stock and Yogo (2005), namely as the IV bias relative to the OLS bias, since $\lim_{\rho^2 \rightarrow 1} \sqrt{d'^{\dagger}_{IV}W^\dagger d^{\dagger}_{IV} / (d'_{OLS}W^\dagger d_{OLS})} = \lim_{\rho^2 \rightarrow 1} B(\rho^2, \mu^2) / B(\rho^2, 0) = \lim_{\rho^2 \rightarrow 1} B(\rho^2, \mu^2)$.

When $\mu^2 = 0$, the modal IV and OLS bias coincide, $B(\rho^2, 0) = |\rho|$, such that the criterion in that case equals the modulus of the canonical correlation between the endogenous regressor Y and the structural error u as $T \rightarrow \infty$. Therefore, the bias criterion can also be interpreted in terms of the canonical correlation between the endogenous regressor Y and the structural error u . For $\mu^2 \rightarrow \infty$, $s^+ \rightarrow \infty$ and the OLS bias induced by the correlation between Y and u is fully eliminated by projecting onto the instrument in the first-stage. The case with $0 < \mu^2 < \infty$ approximates the finite-sample case in which some correlation between $\hat{Y} = Z\hat{\Pi}$ and u remains. The value of $B(\rho^2, \mu^2)$ can be thought of as the extent to which the instrumentation step attenuates the effect of the correlation between Y and u by regressing y on \hat{Y} instead of on Y , in units of the endogeneity parameter.

Although the bias criterion in Definition 1 shares the intuitive interpretation of Montiel Olea and Pflueger (2013) – and, for $\rho^2 \rightarrow 1$, Stock and Yogo (2005) – it is ill-suited for comparing bias across different estimators. While

¹³For $\mu^2 \in (0, \underline{\mu}^2)$, the maximum of B over $\rho^2 \in [0, 1)$ is achieved for a value of ρ^2 close to, but strictly below, one.

it is always the case that $\|d_{IV}^{+VAR}\| \leq \|d_{IV}^{+LP}\|$, the ranking of B between VARs and LPs is ambiguous because $\rho_{VAR}^2 \geq \rho_{LP}^2$ but $R^{LP} \geq R^{VAR}$. This ambiguity does not contradict Proposition 4. Rather, it reflects that the bias criterion, B , employs different norms for LPs and VARs based on their distinct structural systems. Specifically, because the asymptotic OLS covariance in LPs generally exceeds that in VARs, the IRF bias coefficients are standardized into (weakly) larger absolute units than in VARs.

4.2 The Weak Instrument Test

In the local-to-zero asymptotic framework, Σ_{wv} and the auxiliary parameters θ and Θ are consistently estimable under Assumptions 1 and 4. In contrast, Assumption 2 implies that the concentration parameter μ^2 and the vector β are not estimable. Consequently, the true model, defined by the pair $\{\mu^2, \beta\}$, remains unknown to the econometrician and can be any element of the parameter space $\{\mu^2 \in \mathbb{R}_{\geq 0}, \beta \in \mathbb{R}^N\}$.

Our objective is a statistical test of the null hypothesis that the true model, $\{\mu^2, \beta\}$, belongs to the subset of weakly-identified models. This weak instrument set is the set of all models for which the bias criterion in Definition 1 is greater than or equal to a chosen tolerance level, τ . As discussed above, the bias criterion, $B(\rho^2, \mu^2)$, depends only on two unknown scalars, μ^2 and ρ^2 . Using the definition of Σ_{wv} in Assumption 4, the endogeneity parameter, ρ^2 , can be expressed as

$$(13) \quad \rho^2 = \frac{(\sigma_{wv}/\sigma_w^2 - \beta)'(\Sigma_{u|w}/\sigma_w^2)^{-1}(\sigma_{wv}/\sigma_w^2 - \beta)}{1 + (\sigma_{wv}/\sigma_w^2 - \beta)'(\Sigma_{u|w}/\sigma_w^2)^{-1}(\sigma_{wv}/\sigma_w^2 - \beta)},$$

where $\Sigma_{u|w} = \Sigma_v - \sigma_{wv}\sigma'_{wv}/\sigma_w^2$ is an observed positive definite matrix. This alternative expression for the endogeneity parameter makes clear that ρ^2 depends on the unobserved parameter β and on the observed covariance structure Σ_{wv} . Without any additional restrictions on β , knowledge of Σ_{wv} does not impose any restrictions on the possible values of $\rho^2 \in [0, 1)$. Therefore, to determine whether a model belongs to the weak instrument set, it suffices to consider only the scalars ρ^2 and μ^2 . Accordingly, we can characterize a model simply by the pair $\{\mu^2, \rho^2\}$. This leads to the following definition of the set of weakly-identified models:

Definition 2. *The weak instrument set is*

$$\mathbb{B}_\tau = \{\rho^2 \in [0, 1); \mu^2 \in \mathbb{R}_{\geq 0} : B(\rho^2, \mu^2) \geq \tau\}.$$

The next proposition provides a threshold for μ^2 that guarantees inclusion in the weak instrument set but does not depend on ρ^2 :

Proposition 5. *Define $m(\tau) = (R + 1)\frac{(1-\tau)^2}{\tau}$. For $m(\tau) \geq \underline{\mu}^2$, $\{\mu^2, \rho^2\} \in \mathbb{B}_\tau$ if and only if $\mu^2 \leq m(\tau)$.*

Proof. See Appendix B.5. □

Proposition (5) implies that – as long as the tolerance level, τ , is not too large – a test of the null hypothesis that the true model belongs to the weak instrument set can be based on a test of the null that the concentration parameter, μ^2 , does not exceed the threshold value, $m(\tau)$. The requirement that $m(\tau) \geq \underline{\mu}^2 = 2\left(\sqrt{1 + (R + 1)^2} - (R + 1)\right)$ is not restrictive in practice. The condition is always satisfied, for example, for $\tau \leq 0.5$, which well exceeds the level of bias most researchers consider acceptable in practice.¹⁴

Given Proposition 5, the null and alternative hypotheses for the test can be stated formally as:

$$(14) \quad H_0 : \mu^2 \leq m(\tau) \quad , \quad H_1 : \mu^2 > m(\tau).$$

The test can be based on the usual F -statistic,

$$(15) \quad F = \frac{1}{\hat{\sigma}_w^2} \frac{Y'ZZ'Y}{T} = T \left(\frac{\hat{\Pi}}{\hat{\sigma}_w} \right)^2,$$

where $\hat{\sigma}_w = (Y - \hat{\Pi}Z)'(Y - \hat{\Pi}Z)/T$. Under Assumptions 1-4, $T^{\frac{1}{2}}Y'Z/\sigma_w$ converges in distribution to a normal distribution with mean c/σ_w and unit variance. Since $\hat{\sigma}_w \xrightarrow{p} \sigma_w$, it follows that when $\mu^2 = m(\tau)$,

$$(16) \quad F \xrightarrow{d} \chi_1^2(m(\tau)),$$

¹⁴The most common choice in existing tests is $\tau = 0.10$. The threshold $\underline{\mu}^2$ is decreasing in R , and $\underline{\mu}^2 \rightarrow 0$ for $R \rightarrow \infty$, such that the condition can be essentially ignored for large R .

where $\chi_1^2(m(\tau))$ denotes the chi-squared distribution with non-centrality $m(\tau)$ and one degree of freedom. The following Proposition states the asymptotic validity of the test for significance level α :

Proposition 6. *Under Assumptions 1-4 and $m(\tau) \geq \underline{\mu}^2$, the weak instrument test is pointwise asymptotically valid, that is*

$$\sup_{\mu^2 \leq m(\tau)} \lim_{T \rightarrow \infty} \text{Prob}(F > C(m(\tau), \alpha)) = \alpha,$$

where $C(\mu^2, \alpha)$ is the upper α -quantile of $\chi_1^2(\mu^2)$.

Proof. See Appendix B.6. □

Table 1 reports critical values for significance levels $\alpha = 0.01, 0.05,$ and 0.10 and bias tolerance levels $\tau = 0.05, 0.10,$ and 0.20 .¹⁵ The null hypothesis of a weak instrument is rejected whenever the observed F -statistic exceeds the critical value corresponding to the chosen significance and tolerance levels. For the common choices $\alpha = 0.05$ and $\tau = 0.10$, the critical value for an individual IRF coefficient ($R = 1$) is 32.1, substantially larger than the widely used rule-of-thumb value of 10 or the Montiel Olea and Pflueger (2013) critical value of around 23. Moreover, the critical values increase with the rank of the impulse response R , which depends on the IRF estimator as described at the end of Section 2.3. As a result, compared to current testing procedures – which have the shortcomings mentioned above – the test based on the modal IV bias will fail to reject more frequently, and therefore often reach different conclusions regarding the potential for weak instrument bias.

4.3 Extensions

4.3.1 Bootstrap Refinement and Simulations

In finite samples, uncertainty in auxiliary parameters and departures from normality of the regression scores can distort the size of the weak instrument test, as the non-central χ^2 critical values in Table 1 apply only in the asymptotic

¹⁵ Critical values for cases not covered in Table 1 are straightforward to obtain in standard software: `ncx2inv(1- α , 1, $m(\tau)$)` in Matlab, `invnchi2(1, $m(\tau)$, 1- α)` in Stata, or `qchisq(1- α , df=1, ncp= $m(\tau)$)` in R, where $m(\tau) = (R + 1)(1 - \tau)^2/\tau$. Since $\lim_{R \rightarrow \infty} R^{-1}C(m(\tau), \alpha) = (1 - \tau)^2/\tau$, for large R , a useful rule-of-thumb is a critical value of $R(1 - \tau)^2/\tau$.

TABLE 1: ASY. CRITICAL VALUES FOR THE WEAK INSTRUMENT TEST

		$\alpha = 0.01$			$\alpha = 0.05$			$\alpha = 0.10$		
R	τ	0.05	0.10	0.20	0.05	0.10	0.20	0.05	0.10	0.20
1	1	69.5	40.3	23.6	58.6	32.1	17.4	53.1	28.2	14.5
2	2	93.8	52.6	29.4	81.1	43.2	22.5	74.7	38.6	19.2
3	3	117.1	64.3	34.9	102.9	53.8	27.3	95.6	48.6	23.6
4	4	139.9	75.5	40.0	124.2	64.1	31.9	116.2	58.5	27.9
5	5	162.1	86.4	45.0	145.2	74.2	36.3	136.6	68.1	32.1
6	6	184.1	97.1	49.8	166.0	84.2	40.7	156.8	77.6	36.2
7	7	205.7	107.7	54.6	186.6	94.0	45.0	176.8	87.1	40.2
8	8	227.2	118.0	59.2	207.1	103.7	49.2	196.8	96.4	44.2
9	9	248.4	128.3	63.7	227.4	113.3	53.3	216.6	105.7	48.1
10	10	269.5	138.4	68.2	247.6	122.9	57.4	236.3	114.9	52.0
11	11	290.5	148.5	72.6	267.7	132.3	61.5	256.0	124.1	55.9
12	12	311.3	158.5	77.0	287.7	141.8	65.5	275.6	133.2	59.8
13	13	332.1	168.4	81.4	307.7	151.1	69.5	295.1	142.3	63.6
14	14	352.7	178.2	85.6	327.6	160.5	73.5	314.6	151.4	67.4
15	15	373.3	188.0	89.9	347.4	169.8	77.4	334.0	160.4	71.2
16	16	393.8	197.7	94.1	367.2	179.0	81.4	353.4	169.4	74.9
17	17	414.2	207.4	98.3	386.9	188.2	85.3	372.7	178.4	78.7
18	18	434.5	217.0	102.5	406.6	197.4	89.2	392.1	187.3	82.4
19	19	454.8	226.6	106.6	426.2	206.6	93.0	411.3	196.3	86.1
20	20	475.0	236.2	110.8	445.8	215.7	96.9	430.6	205.2	89.9
24	24	555.5	274.1	127.0	523.8	252.0	112.1	507.3	240.6	104.6
28	28	635.3	311.6	143.0	601.4	288.0	127.2	583.7	275.8	119.1
32	32	714.6	348.8	158.8	678.6	323.8	142.1	659.8	310.8	133.6
36	36	793.5	385.7	174.4	755.6	359.4	156.9	735.7	345.7	147.9
40	40	872.0	422.3	189.9	832.2	394.8	171.6	811.4	380.5	162.2
48	48	1028.2	495.0	220.5	985.0	465.1	200.7	962.3	449.6	190.5
60	60	1260.8	602.9	265.6	1212.9	569.9	243.9	1187.7	552.7	232.7
72	72	1492.0	709.8	310.1	1439.8	674.0	286.6	1412.3	655.3	274.4
84	84	1721.9	816.0	354.1	1665.8	777.5	329.0	1636.3	757.4	315.9
96	96	1950.9	921.5	397.8	1891.2	880.6	371.1	1859.7	859.2	357.2
108	108	2179.2	1026.6	441.1	2116.1	983.4	412.9	2082.8	960.7	398.3
120	120	2406.9	1131.2	484.2	2340.5	1085.8	454.6	2305.5	1062.0	439.3

Note: The test rejects the null hypothesis of a weak instrument if the F -statistic in (15) exceeds the critical value. The critical value is a function of the significance level, α ; the rank of the impulse response R ; and the bias tolerance, τ . The critical values shown are based on the non-central χ^2 with non-centrality $m(\tau)$ and one degree of freedom.

limit $T \rightarrow \infty$ (Assumptions 1 and 4). For improvements in performance in finite samples, critical values for the F -statistic can alternatively be obtained via a bootstrap procedure. Appendix C describes a parametric bootstrap and presents simulation evidence demonstrating that the test performs well across a range of DGPs in terms of size and power. The simulation results, which include several DGPs based on empirical SVAR models from the literature, show that the performance of the test under the asymptotic χ_1^2 critical values is reasonably good overall, but the bootstrap procedure effectively corrects remaining small-sample size distortions. The simulations also indicate that the test has meaningful power.

4.3.2 Testing with Parameter Restrictions

In macroeconomic applications, samples are usually short, and the share of plausibly exogenous variation in the normalizing variables is often small. In practice, observed F -statistics are often lower than the threshold of 32.1 for a single IRF coefficient using $\tau = 0.10$ and $\alpha = 0.05$. Moreover, the critical values for joint tests rise with the IRF rank, R , and therefore yield even larger critical values. In LP models, for example, R equals the number of IRF coefficients. Rejecting a bias of $\tau = 0.10$ or more at $\alpha = 0.05$ for an IRF with 24 monthly horizons requires an F -statistic exceeding 252.0 ($R = 24$). Rejecting the same weak instrument null in a four-variable VAR requires an F -statistic above 53.8 ($R = 3$), substantially less than the critical value of 252.0 in the LP model, but still larger than F -statistics typically obtained in empirical macro applications.

Failure to reject the weak instrument null does not imply that the bias of the true DGP necessarily exceeds the tolerance level. It only indicates that the bias exceeds the threshold in the worst-case scenario of near-perfect endogeneity ($\rho^2 \rightarrow 1$). These worst-case models may represent highly unrealistic scenarios in practice, making the weak instrument set in Definition 2 potentially unnecessarily conservative. In fact, to generate the largest possible bias, the true IRF must always contain at least one coefficient that is arbitrarily large in absolute value. To see why, consider the formulation of the endogeneity parameter in (13). The expression shows that ρ^2 is increasing in $(\sigma_{wv}/\sigma_w^2 - \beta)'(\Sigma_{u|w}^{-1}/\sigma_w^2)(\sigma_{wv}/\sigma_w^2 - \beta)$, which has an upper bound at

$\max_{\text{eval}}\{\Sigma_{u|w}^{-1}/\sigma_w^2\}(\sigma_{wv}/\sigma_w^2 - \beta)'(\sigma_{wv}/\sigma_w^2 - \beta)$. Therefore, achieving $\rho^2 \rightarrow 1$ requires $\|\sigma_{wv}/\sigma_w^2 - \beta\| \rightarrow \infty$, which in turn requires $\|\beta\| \rightarrow \infty$. Thus, the worst-case DGPs that determine the boundary of the weak instrument set must feature IRF coefficients that are arbitrarily large in absolute value for an exogenous unit innovation in Y_t .

In applied work, researchers may often be comfortable ruling out DGPs with such extreme dynamic causal effects. Doing so can be done by imposing restrictions on the unknown structural parameters in β . Depending on the specific application, these restrictions can be based on theoretical considerations, additional external information, or simply ex-ante plausibility considerations. To obtain a constrained weak instrument test with greater power, the restrictions must bound the endogeneity parameter ρ^2 (or $\tilde{\rho}^2$) away from one, i.e. $\rho^2 \leq 1 - \varepsilon$ for some $\varepsilon > 0$. In SVAR models, this amounts to imposing a lower bound on the fraction of the variance of Y – the one-step head forecast error of the scaling variable – that is explained by the identified structural shock.

To illustrate how parameter restrictions can limit the set of admissible DGPs and – provided these restrictions indeed hold – achieve greater power, consider the following simple structural AR(1) model as an example:

$$(17) \quad x_t^y = \lambda x_{t-1}^y + e_t, \quad e_t = \delta_0 Y_t + u_t, \quad Y_t = \frac{\mu}{\sqrt{T}} Z_t + w_t$$

$$[e_t, w_t, Z_t] \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, V), \quad V = \begin{bmatrix} 1 & r & 0 \\ r & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix},$$

where the covariance, $0 \leq r \leq 1$, and autoregressive parameter, λ , are consistently estimable, but the impact coefficient δ_0 and first-stage parameter μ are not. Let $\lambda_H = [1, \lambda, \dots, \lambda^{H-1}]'$, so the true impulse response vector is $\delta = \lambda_H \delta_0$. The unobserved endogeneity parameter is $\rho^2 = (r - \delta_0)^2 / ((r - \delta_0)^2 + (1 - r^2))$ and the asymptotic OLS bias is $d_{OLS} = \lambda_H (r - \delta_0)$. The unconstrained worst-case DGP with $\rho^2 \rightarrow 1$ requires making $|\delta_0|$ arbitrarily large. To rule out such extreme impulse responses, consider imposing box constraints, $\underline{\delta} \leq \delta_0 \leq \bar{\delta}$, where $|\underline{\delta}|, |\bar{\delta}| < \infty$. With these constraints, ρ^2 has the sharp bound $\bar{\rho}^2 = \bar{m} / (\bar{m} + 1 - r^2) < 1$, where $\bar{m} = \max\{(r - \underline{\delta})^2, (r - \bar{\delta})^2\}$, such that $B \leq |\bar{\rho}| / (1 + s_1^\dagger)$ for $\mu^2 \geq \underline{\mu}^2$. Given r , $|\bar{\rho}| / (1 + s_1^\dagger) = \tau$ can be solved numerically for the threshold value, $m^c(\tau)$, using the algorithm in Proposition 2

to compute s_1^+ . Provided this threshold exceeds $\underline{\mu}^2$, it can then be used as the non-centrality parameter to generate an asymptotically valid $\chi_1^2(m^c(\tau))$ critical value. Since $\bar{\rho}$ has a sharp bound away from unity, this critical value will be lower than that of the unconstrained test. In finite samples, the covariance r is unknown but can be replaced with a consistent estimate \hat{r} .

For a general formulation of the weak instrument test with parameter constraints, we make the following assumption:

Assumption 5. $\max_{\beta: Q(\beta) \leq 0} \rho^2(\beta) = \bar{\rho}^2 < 1 - \varepsilon$, for some $\varepsilon > 0$.

Here, we have made the dependence of ρ^2 on β in (13) explicit. This assumption states that some constraints $Q(\beta) \leq 0$ are available to the econometrician that are sufficient – given Σ_{wv} – to bound the endogeneity parameter ρ^2 strictly below one. The set of weakly-identified admissible models is defined as:

Definition 3. *The constrained weak instrument set is*

$$\mathbb{B}_\tau^c = \{ \rho^2 \in [0, \bar{\rho}^2]; \mu^2 \in \mathbb{R}_{\geq 0} : B(\rho^2, \mu^2) \geq \tau \}.$$

The following threshold for μ^2 guarantees inclusion in \mathbb{B}_τ^c :

Proposition 7. *Let $m^c(\tau) \in \mathbb{R}_{\geq 0}$ be the unique value of μ^2 implicitly defined by $|\bar{\rho}|/(1 + s_R^+) = \tau$. For $m^c(\tau) \geq \underline{\mu}^2$, $\{\mu^2, \rho^2\} \in \mathbb{B}_\tau^c$ if and only if $\rho^2 \leq \bar{\rho}^2$ and $\mu^2 \leq m^c(\tau)$.*

Proof. See Appendix B.7 □

Since the upper bound $\bar{\rho}^2$ depends on Σ_{wv} , in practice we must use an estimate $\hat{\rho}^2$ based on $\hat{\Sigma}_{wv} = [w \ v]'[w \ v]/T$, which in turns leads to an estimate of the threshold, $\hat{m}^c(\tau)$. Asymptotic validity is established next:

Proposition 8. *Under Assumptions 1-5 and $m^c(\tau) \geq \underline{\mu}^2$, the constrained weak instrument test is pointwise asymptotically valid, that is*

$$\sup_{\mu^2 \leq m^c(\tau)} \lim_{T \rightarrow \infty} \text{Prob}(F > C(\hat{m}^c(\tau), \alpha)) = \alpha,$$

where $C(\mu^2, \alpha)$ is the upper α -quantile of $\chi_1^2(\mu^2)$.

Proof. See Appendix B.8. □

Unlike the unconstrained test, the critical values for the constrained test depend not only on R , but also on the constraints $Q(\beta) \leq 0$ and the sample covariance $\hat{\Sigma}_{wv}$, which jointly determine the estimate of the upper bound on the endogeneity parameter, $\bar{\rho}^2$. Conditional on this bound, the threshold estimate $\hat{m}^c(\tau)$ is computed using the iterative algorithm in Proposition 2. This algorithm is computationally inexpensive and guaranteed to converge to the mode. Once $\hat{m}^c(\tau)$ is established, the critical values are readily obtained using standard statistical software commands.¹⁶

4.3.3 Robust Versions of the Test

The testing procedures above are based on the assumption of conditional homoskedasticity and serially uncorrelated regression scores (CHSU), see Assumption 4. While a common assumption in instrumental variable IRF applications, it may not always be appropriate. To allow for heteroskedasticity or serial correlation, suppose that $[\eta_1, \eta_2]'$ in Assumption 4 follows a normal distribution with some covariance $\Omega_{wv} \neq \Sigma_{wv}$. In that case, Propositions 1-3 continue to hold after replacing Σ_{wu} with Ω_{wu} containing the elements ω_w^2 , $\omega_{wu} = \omega_{wv} - \beta\omega_w^2$, and $\Omega_u = \Omega_v - \beta\beta'\omega_w^2 - \omega_{wu}\beta' - \omega'_{wu}$. What changes, however, is that ω_{wu}/ω_w^2 is no longer the asymptotic bias of the OLS estimate, $\hat{\beta}_{OLS}$. Consequently, $\Theta\omega_{wu}/\omega_w^2$ is also no longer equal to the asymptotic bias of $\hat{\delta}_{OLS}$. We describe two options to accommodate a general covariance Ω_{wv} that differ in the choice of bias criterion. The first option maintains the Montiel Olea and Pflueger (2013) bias criterion, in which case researchers can simply proceed as in the CHSU case but using the robust F -statistic. The second option maintains the asymptotic OLS bias and covariance in the criterion, which results in application-specific critical values for the robust F -statistic that depend on consistent estimates of Ω_{wv} and $\Sigma_{Yy} = E[[y, Y]'[y, Y]]$ and must be obtained numerically.

Option 1: Bias criterion of Montiel Olea and Pflueger (2013):

$$(18) \quad B_1^{het} = \sqrt{\frac{d_{IV}^{+'} W_{\Omega}^{\dagger} d_{IV}^{+}}{1 + d' W_{\Omega}^{\dagger} d}},$$

¹⁶These are as in footnote 15, but replacing $m(\tau)$ with $\hat{m}^c(\tau)$.

where $d = \Theta\omega_{wu}/\omega_w^2$, d_{IV}^{+} is defined as in Proposition 3 after replacing d_{OLS} with d and using $\rho^2 = \omega'_{wu}\Omega_u^{-1}\omega_{wu}/\omega_w^2$, and W_Ω is defined as in Definition 1 but using the Cauchy scale matrix $\Omega_{u|w}/\omega_w^2$ implied by Ω_{wu} . As a benchmark, this criterion maintains the modal IV bias in the most adverse scenario of an irrelevant instrument and perfect endogeneity. As in the CHSU case, the bias criterion simplifies to $0 \leq B_1^{het} = |\rho|/(1 + s_R^+) < 1$, where s_R^+ depends on R , μ^2 and ρ^2 . The test can be based on the robust F -statistic $T(\hat{\Pi}/\hat{\sigma}_w^{rob})^2$, where $\hat{\sigma}_w^{rob}$ is a heteroskedasticity-autocorrelation robust estimate of the standard deviation of the residual in the first-stage equation, and the critical values in Table 1 remain applicable.

Option 2: Bias criterion based on the OLS benchmark:

$$(19) \quad B_2^{het} = \sqrt{\frac{\bar{d}_{IV}^{+}W^\dagger\bar{d}_{IV}^{+}}{1 + d'_{OLS}W^\dagger d_{OLS}}}$$

where W^\dagger and d_{OLS} are defined using $\sigma_w^2 = \sigma_Y^2$, $\sigma_{wu} = \sigma_{Yy}/\sigma_Y^2 - \beta$ and $\Sigma_u = \Sigma_y - \beta\beta'\sigma_w^2 - \sigma_{wu}\beta' - \beta\sigma'_{wu}$, and $\Sigma_{Yy} = \begin{bmatrix} \sigma_Y^2 & \sigma'_{Yy} \\ \sigma_{Yy} & \Sigma_y \end{bmatrix}$. The bias criterion can be written as

$$(20) \quad B_2^{het} = \sqrt{\frac{1}{1 + s_R^+}} \cdot \sqrt{\frac{(\omega_{wv}/\omega_w^2 - \beta)'W^\dagger(\omega_{wv}/\omega_w^2 - \beta)}{1 + (\sigma_{Yy}/\sigma_Y^2 - \beta)'W^\dagger(\sigma_{Yy}/\sigma_Y^2 - \beta)}}$$

where s_R^+ depends on R , μ^2 , and $\rho^2 = \frac{(\omega_{wv}/\omega_w^2 - \beta)'W_\Omega^\dagger(\omega_{wv}/\omega_w^2 - \beta)}{1 + (\omega_{wv}/\omega_w^2 - \beta)'W_\Omega^\dagger(\omega_{wv}/\omega_w^2 - \beta)}$. The supremum of $B_2^{het} \geq 0$ is generally no longer characterized by $\rho^2 \rightarrow 1$ and the bias criterion can exceed unity. In practice, obtaining the supremum as a function of μ^2 requires numerical optimization over β , and evaluating (20) requires estimates of Ω_{wv} and Σ_{Yy} .¹⁷ The critical values in Table 1 are not applicable, but need to be computed separately for every application. The relevant test statistic is the robust F -statistic.

¹⁷The optimization problem can feature local maxima. We obtained good results using the value that maximizes the second term in (20) as a starting value in standard optimization algorithms. This value is given by $\frac{\sigma_{Yy}}{\sigma_Y^2} + \frac{(\sigma_{Yy}/\sigma_Y^2 - \omega_{wv}/\omega_w^2)}{(\omega_{wv}/\omega_w^2 - \sigma_{Yy}/\sigma_Y^2)'W^\dagger(\omega_{wv}/\omega_w^2 - \sigma_{Yy}/\sigma_Y^2)}$.

5 Weak Instrument Bias in Empirical Applications

A comprehensive empirical assessment of weak instrument bias in applications of IV-based IRF estimators is beyond the scope of this paper. Instead we examine four structural VAR studies that are broadly representative for the range of scenarios researchers are likely to encounter when applying our test:

1. Monetary Policy Shocks (MP): Based on Gertler and Karadi (2015), this application uses a monthly VAR(12) in the 1-year Treasury rate, industrial production, the CPI, and the excess bond premium. The scaling variable is the 1-year Treasury rate, and the shocks are identified using high-frequency changes in interest rate futures around FOMC announcements.
2. Oil Supply Shocks (OIL): Following Montiel Olea et al. (2021), we estimate a monthly VAR(12) in oil production, the Kilian (2009) index of global economic activity, and the real price of oil. Oil production serves as the scaling variable, and shocks are identified using a monthly version of the Kilian (2008) exogenous oil supply shock measure.
3. Uncertainty Shocks (UNC): Drawn from Carriero et al. (2015), this model is a monthly VAR(12) comprising the VXO volatility index, industrial production, employment, hours worked, the CPI, real wages, the Federal Funds Rate, and stock prices. The scaling variable is the VXO, and the shock is identified with the geopolitical event indicator constructed by Bloom (2009).
4. Tax Policy Shocks (TAX): Based on Mertens and Montiel Olea (2018), this annual VAR(1) includes average marginal tax rates, average adjusted gross income, real GDP, unemployment, inflation, the federal funds rate, government spending, debt, and stock prices. The scaling variable is the marginal tax rate, and the shocks are identified using a narrative series of historical tax reforms as the instrument.

In each application, the instrumental variable is the residual in the regression of the identifying series on the lag sequence in the VAR system.

The first column in Table 2 reports the sample F -statistic in (15) for each application. All F -statistics exceed 10, such that researchers would reject the null of a weak instrument based on the common rule-of-thumb associated with

$\alpha = 0.05, \tau = 0.10$. The F -statistics for the OIL, UNC and TAX applications also exceed the Montiel Olea and Pflueger (2013) threshold of 23.

The subsequent four columns in Table 2 provide bootstrap critical values for the baseline test (Section 4) at significance levels $\alpha \in \{0.05, 0.10\}$ and bias thresholds $\tau \in \{0.10, 0.20\}$. Panel (a) reports critical values for pointwise tests of individual IRF coefficients. Notably, our proposed test yields results that frequently diverge from existing procedures. While the pointwise test consistently rejects the weak instrument null in the UNC and TAX applications, results for the OIL application reject only when $\tau = 0.20$. In addition, the test never rejects the null in the MP application. Panel (b) reports critical values for the simultaneous test, which considers the IRF coefficients of all VAR variables jointly across any combination of forecast horizons. While this joint test consistently rejects the null of weak identification in the UNC application, it fails to reject in the remaining three applications.

In conclusion, while weak instrument bias appears to be a negligible concern for the UNC application, the evidence for the OIL and TAX applications remains mixed at best. For the MP application, the test results provide no basis to rule out substantial weak instrument bias.

The unconstrained test is designed to be robust against the worst-case DGP, characterized by the limit $\rho^2 \rightarrow 1$. As established above, this condition requires that at least one IRF coefficient under consideration becomes arbitrarily large in absolute value. The condition also requires the variance contribution of the targeted structural shock to the innovations in the scaling variable innovations to vanish (i.e., $1 - \rho^2 \rightarrow 0$). In the MP application, for instance, the worst-case scenario represents a DGP where exogenous monetary policy shocks play no role for interest rate dynamics; in the OIL application, the worst-case occurs when oil supply shocks play no role in driving oil production, and so on.

Researchers may choose to limit the set of admissible DGPs by imposing additional structural restrictions, $Q(\beta) \leq 0$, to exclude these extreme scenarios. As discussed in Section 4.3.2, for such restrictions to meaningfully affect the set of weakly-identified models, they must bound the endogeneity parameter ρ^2 away from unity. This is equivalent to enforcing a lower bound on the structural shock's variance contribution, $1 - \rho^2$. While various constraints – such as sign or magnitude restrictions on structural parameters or IRFs – can

TABLE 2: TEST RESULTS IN EMPIRICAL APPLICATIONS

		No Constraints on DGP: Bootstrap Critical Value				With Constraints on DGP: Minimum Required $1 - \rho^2$			
		$\tau = 0.10$		$\tau = 0.20$		$\tau = 0.10$		$\tau = 0.20$	
F -stat		$\alpha = 0.05$	0.10	0.05	0.10	$\alpha = 0.05$	0.10	0.05	0.10
<i>(a) Pointwise Test</i>									
MP	10.8	33.9	29.2	18.2	15.0	0.93	0.89	0.71	0.53
OIL	29.3	35.2	30.1	19.5	15.7	0.38	0.06	–	–
UNC	255.2	34.5	30.2	19.3	15.9	–	–	–	–
TAX	43.6	42.5	34.9	22.6	17.9	–	–	–	–
<i>(b) Joint Test</i>									
MP	10.8	56.1	50.5	29.2	25.1	0.96	0.94	0.84	0.77
OIL	29.3	46.3	41.3	24.2	20.8	0.65	0.53	–	–
UNC	255.2	99.2	91.0	48.1	42.6	–	–	–	–
TAX	43.6	110.5	99.3	53.7	45.3	0.83	0.78	0.33	0.07

Notes: Test results for the weak instrument null for individual IRF coefficients (Panel a) and for all IRF coefficients jointly (Panel b). The last four columns show the minimum required variance contribution of the targeted shock (the smallest $1 - \rho^2$) that parameter restrictions must imply in order to reject the weak instrument null. Entries marked “–” indicate that the null is rejected for all DGPs, without requiring any parameter restrictions.

generate such a bound, their plausibility is inherently application-specific.

Rather than evaluating specific restrictions for each application, we pose an alternative question: What is the minimum variance contribution, $1 - \rho^2$, required to reject the weak instrument null? A required value near one implies that the test only rejects if the true DGP yields nearly unbiased OLS estimates; in a VAR context, this means the true IRF must closely resemble one generated by a Cholesky decomposition with the scaling variable ordered first. Conversely, a relatively small required value for $1 - \rho^2$ suggests that even modest parameter restrictions may be sufficient to alter the test’s outcome. The minimum required variance contribution thus provides a general metric for quantifying the degree of restrictiveness required across applications.

For cases where the unconstrained test fails to reject weak identification, the final four columns of Table 2 report the minimum variance contribution,

$1 - \rho^2$, that is required for a constrained test to reject the null. In the OIL application, relatively small variance contributions of supply shocks to monthly oil production fluctuations are sufficient to reject in pointwise tests; however, the joint test requires a contribution of at least one half. In the TAX application, modest variance contributions from tax shocks allow for rejection in the joint test at $\tau = 0.20$, though substantially higher values are required when $\tau = 0.10$. The results for the MP application are relatively clear: To reject the null at $\tau = 0.20$ and $\alpha = 0.10$ in a pointwise test, monetary policy shocks must account for at least one-half of the variance of monthly interest rate innovations. At the conventional levels of $\tau = 0.10$ and $\alpha = 0.05$, monetary policy shocks must explain nearly all variation (93 percent) in the interest rate to reject weak identification. In the joint test, the monetary policy shocks must explain a large majority of the monthly innovations in the interest rate for any of the values of τ and α reported in Table 2. Consequently, ruling out weak identification in the MP application necessitates the assumption of a DGP in which a recursive identification scheme yields approximately unbiased impulse responses. We conclude that the high-frequency instrument used in the Gertler and Karadi (2015) VAR does not appear to add much useful information for reliably identifying the effects of monetary policy shocks on macroeconomic aggregates.

6 Concluding Remarks

This paper provides theoretical insights and a novel toolkit for addressing weak instrument bias in just-identified IV-based impulse response estimators. While developed in the context of IRFs, these results extend naturally to general IV estimation with multiple outcome variables and to the canonical just-identified single-equation model. Looking ahead, our distributional results offer a foundation for developing size-based weak instrument tests or adjusted critical values for t -inference, following the approach of Lee et al. (2022). Further research could expand these methods to over-identified IRF models or specifications involving multiple endogenous variables, such as those in Mertens and Ravn (2013) and Mertens and Montiel Olea (2018). Additionally, our analytical findings highlight potential pitfalls in simulation-based assessments of IRF estimators (Li et al. 2024); specifically, moment-based criteria like mean-

squared error may be ill-defined in the context of weak identification. Finally, future studies might explore alternative asymptotic frameworks and incorporate other sources of finite-sample bias, such as lag truncation (Montiel Olea et al. 2024).

References

- Andrews, Isaiah, James H. Stock, and Liyang Sun (2019). “Weak Instruments in Instrumental Variables Regression: Theory and Practice”. In: *Annual Review of Economics* 11, pp. 727–753.
- Angelini, Giovanni, Giuseppe Cavaliere, and Luca Fanelli (2024). “An identification and testing strategy for proxy-SVARs with weak proxies”. In: *Journal of Econometrics* 238.2, p. 105604.
- Angrist, Joshua D. and Michal Kolesár (2024). “One Instrument to Rule Them All: The Bias and Coverage of Just-ID IV”. In: *Journal of Econometrics* 240.2.
- Bloom, Nicholas (2009). “The Impact of Uncertainty Shocks”. In: *Econometrica* 77.3, pp. 623–685.
- Carriero, Andrea, Haroon Mumtaz, Konstantinos Theodoridis, and Angeliki Theophilopoulou (2015). “The Impact of Uncertainty Shocks under Measurement Error: A Proxy SVAR Approach”. In: *Journal of Money, Credit and Banking* 47.6, pp. 1223–1238.
- Fieller, E. C. (1932). “The Distribution of the Index in a Normal Bivariate Population”. In: *Biometrika* 24.3/4, pp. 428–440.
- Gertler, Mark and Peter Karadi (2015). “Monetary Policy Surprises, Credit Costs, and Economic Activity”. In: *American Economic Journal: Macroeconomics* 7.1, pp. 44–76.
- Giannone, Domenico, Michele Lenza, and Giorgio E. Primiceri (2025). “Bayesian Inference in IV Regressions”. In: Unpublished manuscript.
- Hausman, Jerry A. and William E. Taylor (1983). “Identification in Linear Simultaneous Equations Models with Covariance Restrictions: An Instrumental Variables Interpretation”. In: *Econometrica* 51.5, pp. 1527–1549.
- Herbst, Edward P and Benjamin K Johannsen (2024). “Bias in local projections”. In: *Journal of Econometrics* 240.1, p. 105655.

- Hinkley, D. V. (1969). “On the Ratio of Two Correlated Normal Random Variables”. In: *Biometrika* 56.3, pp. 635–639.
- Jentsch, Carsten and Kurt G. Lunsford (Oct. 2022). “Asymptotically Valid Bootstrap Inference for Proxy SVARs”. In: *Journal of Business & Economic Statistics* 40.4, pp. 1876–1891.
- Kilian, Lutz (2008). “Exogenous Oil Supply Shocks: How Big Are They and How Much Do They Matter for the U.S. Economy?” In: *Review of Economics and Statistics* 90.2, pp. 216–240.
- (2009). “Not All Oil Price Shocks Are Alike: Disentangling Demand and Supply Shocks in the Crude Oil Market”. In: *American Economic Review* 99.3, pp. 1053–1069.
- Kilian, Lutz and Yun Jung Kim (2011). “How Reliable Are Local Projection Estimators of Impulse Responses?” In: *The Review of Economics and Statistics* 93.4, pp. 1460–1466.
- Kleibergen, Frank and Eric Zivot (May 2003). “Bayesian and Classical Approaches to Instrumental Variable Regression”. In: *Journal of Econometrics* 114.1, pp. 29–72.
- Lee, David S., Justin McCrary, Marcelo J. Moreira, and Jack Porter (2022). “Valid t-Ratio Inference for IV”. In: *American Economic Review* 112.10, pp. 3260–90.
- Lewis, Daniel J. and Karel Mertens (2025). “A Robust Test for Weak Instruments for 2SLS with Multiple Endogenous Regressors”. In: *Review of Economic Studies*.
- Li, Dake, Mikkel Plagborg-Møller, and Christian K. Wolf (2024). “Local projections vs. VARs: Lessons from thousands of DGPs”. In: *Journal of Econometrics*. Forthcoming.
- Ludwig, Julian F. (2024). *Local Projections Are VAR Predictions of Increasing Order*. Tech. rep. SSRN.
- Lunsford, Kurt G. (2016). *Identifying Structural VARs with a Proxy Variable and a Test for a Weak Proxy*. Working Paper 16-05. Federal Reserve Bank of Cleveland Working Paper. Federal Reserve Bank of Cleveland.
- Mariano, R. S. and J. B. McDonald (1979). “A Note on the Distribution Functions of LIML and 2SLS Structural Coefficient in the Exactly Identified Case”. In: *Journal of the American Statistical Association* 74, pp. 847–848.

- Marsaglia, George (1965). “Ratios of Normal Variables and Ratios of Sums of Uniform Variables”. In: *Journal of the American Statistical Association* 60, pp. 193–204.
- Mertens, Karel and José Luis Montiel Olea (2018). “Marginal Tax Rates and Income: New Time Series Evidence”. In: *Quarterly Journal of Economics* 133.4, pp. 1803–1884.
- Mertens, Karel and Morten O. Ravn (2013). “The Dynamic Effects of Personal and Corporate Income Tax Changes in the United States”. In: *American Economic Review* 103.4, pp. 1212–47.
- Montiel Olea, José Luis and Carolin Pflueger (2013). “A Robust Test for Weak Instruments”. In: *Journal of Business & Economic Statistics* 31.3, pp. 358–369.
- Montiel Olea, José Luis, Mikkel Plagborg-Møller, Eric Qian, and Christian K. Wolf (2024). “Double Robustness of Local Projections and Some Unpleasant VARithmetic”. In: arXiv preprint arXiv:2405.09509.
- (2025). “Local Projections or VARs? A Primer for Macroeconomists”. In: *NBER Macroeconomics Annual 2025, Volume 40*. National Bureau of Economic Research.
- Montiel Olea, José Luis, James H. Stock, and Mark W. Watson (2021). “Inference in Structural Vector Autoregressions Identified with an External Instrument”. In: *Journal of Econometrics* 225.1, pp. 74–87.
- Negro, Luigi (2024). “Sample distribution theory using Coarea Formula”. In: *Communications in Statistics - Theory and Methods* 53.5, pp. 1864–1889.
- Nelson, Charles R. and Richard Startz (1990). “Some Further Results on the Exact Small Sample Properties of the Instrumental Variable Estimator”. In: *Econometrica* 58.4, pp. 967–976.
- Plagborg-Møller, Mikkel and Christian K. Wolf (2021). “Local Projections and VARs Estimate the Same Impulse Responses”. In: *Econometrica* 89.2, pp. 955–980.
- Schorfheide, Frank (2005). “VAR forecasting under misspecification”. In: *Journal of Econometrics* 128.1, pp. 99–136.
- Shapiro, Matthew D. and Mark W. Watson (1988). “Sources of Business Cycle Fluctuations”. In: *NBER Macroeconomics Annual* 3, pp. 111–148.
- Staiger, Douglas and James H. Stock (1997). “Instrumental Variables Regression with Weak Instruments”. In: *Econometrica* 65.3, pp. 557–586.

- Stock, James H. and Motohiro Yogo (2005). “Testing for Weak Instruments in Linear IV Regression”. In: *Identification and Inference for Econometric Models*. Ed. by Donald W.K. Andrews. New York: Cambridge University Press, pp. 80–108.
- Yang, Sheng and Zhengtao Gui (2023). “An introduction on the multivariate normal-ratio distribution”. In: *arXiv preprint arXiv:2310.14306*.

APPENDIX

A New Keynesian DGP of Illustrative Example of Section 1.1

The data is simulated from the rational expectations solution to the model:

$$\begin{aligned}
 (A.1) \quad R_t &= \phi_\pi \pi_t + s_t^R, \\
 \text{gap}_t &= E_t \text{gap}_{t+1} - (R_t - E_t \pi_{t+1}) + s_t^d, \\
 \pi_t &= \kappa \text{gap}_t + \beta_d E_t \pi_{t+1} + s_t^s, \\
 s_t^R &= \rho_R s_{t-1}^R + \epsilon_t^R, \\
 s_t^d &= \rho_d s_{t-1}^d + \epsilon_t^d, \\
 s_t^s &= \rho_s s_{t-1}^s + \epsilon_t^s.
 \end{aligned}$$

and $[\epsilon_t^R, \epsilon_t^d, \epsilon_t^s]' \sim \mathcal{N}(\mathbf{0}, \mathcal{I})$. The VAR(1) representation is

$$(A.2) \quad \begin{bmatrix} R_t \\ \text{gap}_t \\ \pi_t \end{bmatrix} = D \begin{bmatrix} \rho_R & 0 & 0 \\ 0 & \rho_d & 0 \\ 0 & 0 & \rho_s \end{bmatrix} D^{-1} \begin{bmatrix} R_{t-1} \\ \text{gap}_{t-1} \\ \pi_{t-1} \end{bmatrix} + D \begin{bmatrix} \epsilon_t^R \\ \epsilon_t^d \\ \epsilon_t^s \end{bmatrix},$$

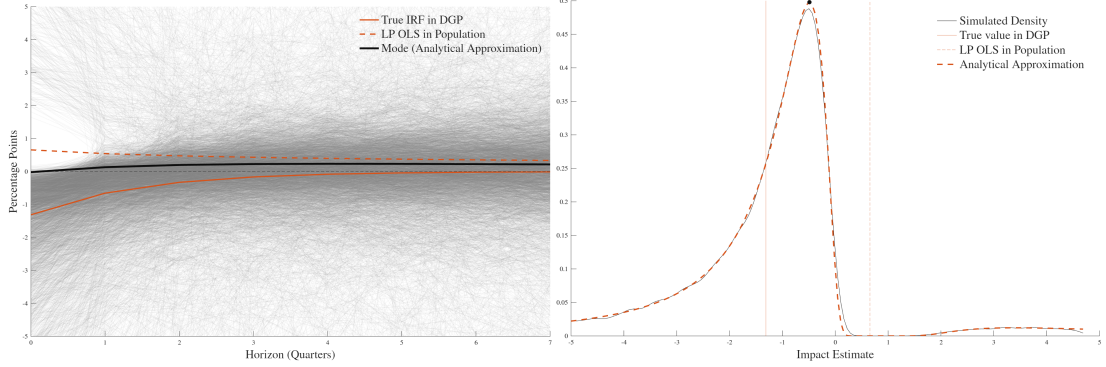
where $D = \begin{bmatrix} 1 - \frac{\frac{\kappa \phi_\pi}{1 - \beta_d \rho_r}}{(1 - \rho_r) + \frac{(\phi_\pi - \rho_r) \kappa}{1 - \beta_d \rho_r}} & \frac{\frac{\kappa \phi_\pi}{1 - \beta_d \rho_d}}{(1 - \rho_d) + \frac{(\phi_\pi - \rho_d) \kappa}{1 - \beta_d \rho_d}} & \frac{\phi_\pi}{1 - \beta_d \rho_s + \frac{\kappa(\phi_\pi - \rho_s)}{1 - \rho_s}} \\ -\frac{1}{(1 - \rho_r) + \frac{(\phi_\pi - \rho_r) \kappa}{1 - \beta_d \rho_r}} & \frac{1}{(1 - \rho_d) + \frac{(\phi_\pi - \rho_d) \kappa}{1 - \beta_d \rho_d}} & -\frac{\frac{\phi_\pi - \rho_s}{1 - \rho_s}}{1 - \beta_d \rho_s + \frac{\kappa(\phi_\pi - \rho_s)}{1 - \rho_s}} \\ -\frac{\frac{\kappa}{1 - \beta_d \rho_r}}{(1 - \rho_r) + \frac{(\phi_\pi - \rho_r) \kappa}{1 - \beta_d \rho_r}} & \frac{\frac{\kappa}{1 - \beta_d \rho_d}}{(1 - \rho_d) + \frac{(\phi_\pi - \rho_d) \kappa}{1 - \beta_d \rho_d}} & \frac{1}{1 - \beta_d \rho_s + \frac{\kappa(\phi_\pi - \rho_s)}{1 - \rho_s}} \end{bmatrix}.$

We set $\phi_\pi = 1.5$, $\kappa = 0.20$, $\beta_d = 0.99$, $\rho_R = 0.5$, $\rho_d = 0.95$, $\rho_s = 0.5$, and generate z_t from

$$(A.3) \quad z_t = \lambda \epsilon_t^R + \sqrt{(1 - \lambda^2)} \epsilon_t^z; \quad E[\epsilon_t^i \epsilon_t^z] = 0, \quad i = R, d, s; \quad \epsilon_t^z \sim \mathcal{N}(0, 1),$$

using $\lambda = 0.75$.

FIGURE A.1: LP-IV Impulse Responses of Inflation to an Interest Rate Shock



Notes: Estimated inflation responses to a one percentage point positive interest rate shock in Monte Carlo samples ($T = 250$) from a simple NK model with Gaussian shocks; see Appendix A. *Left:* LP-IV estimates with a valid instrument and one lag of the interest rate, output gap and inflation as controls. ‘LP OLS in Population’ is the population OLS estimate of the IRF; ‘Mode (Analytical Approximation)’ is the modal LP-IV response from Proposition 3. *Right:* ‘Simulated Density’ is a kernel density estimate of the impact based on 50,000 Monte Carlo samples. Analytical density and major mode (black circle) based on Proposition 3.

B Proofs

B.1 Proof of Proposition 1

Proof. Denoting the joint normal density of the vector $[\nu_1, \nu_2]'$ by $\varphi_\nu(n_1, n_2)$ and using the method of transformations, the density of ν_2/ν_1 is

$$(B.1) \quad \begin{aligned} \varphi(b) &= \int_{-\infty}^{\infty} \det(J) \cdot \varphi_\nu(m, m \cdot b) dm, \\ &= \int_{-\infty}^{\infty} \frac{|m^N| \exp\left(-\frac{1}{2} \begin{bmatrix} m - c \\ m \cdot b \end{bmatrix}' \Sigma_{wu}^{-1} \begin{bmatrix} m - c \\ m \cdot b \end{bmatrix}\right)}{(2\pi)^{\frac{N+1}{2}} \det(\Sigma_{wu})^{\frac{1}{2}}} dm, \end{aligned}$$

where $m = n_1$ and $J = m^N$ is the Jacobian of the inverse transformation $[n_1, n_2] = [m, m \cdot b]'$ with respect to $[m, b]'$. Based on the inverse of a partitioned matrix,

$$(B.2) \quad \Sigma_{wu}^{-1} = \frac{1}{\sigma_w^2(1 - \rho^2)} \begin{bmatrix} 1 & -\sigma'_{wu} \Sigma_u^{-1} \\ -\Sigma_u^{-1} \sigma_{wu} & \sigma_w^2(1 - \rho^2) \Sigma_u^{-1} + \Sigma_u^{-1} \sigma_{wu} \sigma'_{wu} \Sigma_u^{-1} \end{bmatrix},$$

such that

$$(B.3) \quad \begin{bmatrix} m - c \\ m \cdot b \end{bmatrix}' \Sigma_{wu}^{-1} \begin{bmatrix} m - c \\ m \cdot b \end{bmatrix} = \frac{\mu^2}{1 - \rho^2} + d_1(b)m + d_2(b)m^2,$$

where

$$d_1(b) = \frac{2c(b'\Sigma_u^{-1}\sigma_{wu} - 1)}{(1 - \rho^2)\sigma_w^2}; \quad d_2(b) = \frac{(b'\Sigma_u^{-1}\sigma_{wu} - 1)^2 + \sigma_w^2(1 - \rho^2)b'\Sigma_u^{-1}b}{(1 - \rho^2)\sigma_w^2},$$

with $d_2(b) > 0$ since $0 \leq \rho^2 < 1$. Using $\det(\Sigma_{wu}) = \sigma_w^2(1 - \rho^2)\det(\Sigma_u)$ and (B.3),

$$(B.4) \quad \varphi(b) = \int_{-\infty}^{\infty} \det(m^N) \cdot d_0 \cdot \exp\left(-\frac{1}{2}(d_1(b)m + d_2(b)m^2)\right) dm,$$

where $d_0 = \sigma_w^{-1}(1 - \rho^2)^{-\frac{1}{2}}(2\pi)^{-\frac{N+1}{2}}\det(\Sigma_u)^{-\frac{1}{2}}e^{-\frac{1}{2}\frac{\mu^2}{1-\rho^2}}$. Splitting the integral,

$$(B.5) \quad \varphi(b) = d_0 \left[\underbrace{\int_0^{\infty} m^N \exp\left(-\frac{1}{2}(d_1(b)m + d_2(b)m^2)\right) dm}_{f_1(b)} + (-1)^N \underbrace{\int_{-\infty}^0 m^N \exp\left(-\frac{1}{2}(d_1(b)m + d_2(b)m^2)\right) dm}_{f_2(b)} \right].$$

Both integrals always converge since $d_2(b) > 0$. We first solve the integral in $f_1(b)$. Completing the square,

$$(B.6) \quad f_1(b) = e^{\frac{d_1(b)^2}{8d_2(b)}} \int_0^{\infty} m^N \exp\left(-\frac{1}{2}d_2(b)\left(m + \frac{d_1(b)}{2d_2(b)}\right)^2\right) dm.$$

Using the change of variables $m = n - \frac{d_1(b)}{2d_2(b)}$,

$$(B.7) \quad f_1(b) = e^{\frac{d_1(b)^2}{8d_2(b)}} \int_{\frac{d_1(b)}{2d_2(b)}}^{\infty} \left(n - \frac{d_1(b)}{2d_2(b)}\right)^N \exp\left(-\frac{1}{2}d_2(b)n^2\right) dn.$$

Using the binomial expansion,

(B.8)

$$f_1(b) = e^{\frac{d_1(b)^2}{8d_2(b)}} \sum_{j=0}^N \binom{N}{j} \left(-\frac{d_1(b)}{2d_2(b)}\right)^{N-j} \underbrace{\int_{\frac{d_1(b)}{2d_2(b)}}^{\infty} n^j \exp\left(-\frac{1}{2}d_2(b)n^2\right) dn}_{M_j}.$$

The M_j integrals in the summation can be expressed in terms of the gamma function $\Gamma(s) = \gamma(s, 0)$, the lower incomplete gamma function $\gamma(s, v)$, and the upper incomplete gamma function $\Gamma(s, v) = \Gamma(s) - \gamma(s, v)$. Let $q = d_2(b)n^2/2$. If $d_1(b) \geq 0$,

(B.9)

$$M_j = \frac{1}{2} \left(\frac{d_2(b)}{2}\right)^{-\frac{j+1}{2}} \int_{\frac{d_1(b)^2}{8d_2(b)}}^{\infty} q^{\frac{j-1}{2}} e^{-q} dq = \frac{1}{2} \left(\frac{d_2(b)}{2}\right)^{-\frac{j+1}{2}} \Gamma\left(\frac{j+1}{2}, \frac{d_1(b)^2}{8d_2(b)}\right).$$

If instead $d_1(b) < 0$,

(B.10)

$$\begin{aligned} M_j &= \int_0^{\infty} n^j \exp\left(-\frac{1}{2}d_2(b)n^2\right) dn + \int_{\frac{d_1(b)}{2d_2(b)}}^0 n^j \exp\left(-\frac{1}{2}d_2(b)n^2\right) dn, \\ &= \int_0^{\infty} n^j \exp\left(-\frac{1}{2}d_2(b)n^2\right) dn + \int_0^{-\frac{d_1(b)}{2d_2(b)}} (-1)^j (-n)^j \exp\left(-\frac{d_2(b)n^2}{2}\right) d(-n), \\ &= \frac{1}{2} \left(\frac{d_2(b)}{2}\right)^{-\frac{j+1}{2}} \left(\Gamma\left(\frac{j+1}{2}\right) + (-1)^j \int_0^{\frac{d_1(b)^2}{8d_2(b)}} q^{\frac{j-1}{2}} \exp(-q) dq \right), \\ &= \begin{cases} \left(\frac{d_2(b)}{2}\right)^{-\frac{j+1}{2}} \frac{1}{2} \Gamma\left(\frac{j+1}{2}, \frac{d_1(b)^2}{8d_2(b)}\right) & \text{if } j \text{ is odd} \\ \left(\frac{d_2(b)}{2}\right)^{-\frac{j+1}{2}} \left(\Gamma\left(\frac{j+1}{2}\right) - \frac{1}{2} \Gamma\left(\frac{j+1}{2}, \frac{d_1(b)^2}{8d_2(b)}\right) \right) & \text{if } j \text{ is even} \end{cases}. \end{aligned}$$

Following analogous steps for $f_2(b)$ as those leading to (B.8),

(B.11)

$$f_2(b) = e^{\frac{d_1(b)^2}{8d_2(b)}} \sum_{j=0}^N \binom{N}{j} \left(-\frac{d_1(b)}{2d_2(b)}\right)^{N-j} \underbrace{\int_{-\infty}^{\frac{d_1(b)}{2d_2(b)}} n^j \exp\left(-\frac{1}{2}d_2(b)n^2\right) dn}_{N_j}.$$

If $d_1(b) \geq 0$,

$$(B.12) \quad N_j = \begin{cases} -\left(\frac{d_2(b)}{2}\right)^{-\frac{j+1}{2}} \frac{1}{2} \Gamma\left(\frac{j+1}{2}, \frac{d_1(b)^2}{8d_2(b)}\right) & \text{if } j \text{ is odd} \\ \left(\frac{d_2(b)}{2}\right)^{-\frac{j+1}{2}} \left(\Gamma\left(\frac{j+1}{2}\right) - \frac{1}{2} \Gamma\left(\frac{j+1}{2}, \frac{d_1(b)^2}{8d_2(b)}\right)\right) & \text{if } j \text{ is even} \end{cases},$$

and if $d_1(b) < 0$,

$$(B.13) \quad N_j = (-1)^j \left(\frac{d_2(b)}{2}\right)^{-\frac{j+1}{2}} \frac{1}{2} \Gamma\left(\frac{j+1}{2}, \frac{d_1(b)^2}{8d_2(b)}\right).$$

Regardless of the sign of $d_1(b)$, when N is even

$$(B.14) \quad M_j + N_j = \begin{cases} 0 & \text{if } j \text{ is odd} \\ \left(\frac{d_2(b)}{2}\right)^{-\frac{j+1}{2}} \Gamma\left(\frac{j+1}{2}\right) & \text{if } j \text{ is even} \end{cases},$$

such that when N is even,

$$(B.15) \quad \varphi(b) = d_0 e^{\frac{d_1(b)^2}{8d_2(b)}} \sum_{k=0}^{\frac{N}{2}} \binom{N}{2k} \left(-\frac{d_1(b)}{2d_2(b)}\right)^{N-2k} \left(\frac{d_2(b)}{2}\right)^{-\frac{2k+1}{2}} \Gamma\left(\frac{2k+1}{2}\right).$$

When N is odd:

$$(B.16) \quad M_j - N_j = \begin{cases} \left(\frac{d_2(b)}{2}\right)^{-\frac{j+1}{2}} \Gamma\left(\frac{j+1}{2}, \frac{d_1(b)^2}{8d_2(b)}\right) & \text{if } j \text{ is odd} \\ -\text{sgn}(d_1(b)) \left(\frac{d_2(b)}{2}\right)^{-\frac{j+1}{2}} \gamma\left(\frac{j+1}{2}, \frac{d_1(b)^2}{8d_2(b)}\right) & \text{if } j \text{ is even} \end{cases}.$$

Using $\Gamma(s, t) = \Gamma(s) - \gamma(s, t)$ and $-\text{sgn}(t)(-t)^{N-j} = -(-|t|)^{N-j}$,

$$(B.17) \quad \varphi(b) = d_0 e^{\frac{d_1(b)^2}{8d_2(b)}} \left(\sum_{j=1,3,5,\dots}^N \binom{N}{j} \left(-\frac{d_1(b)}{2d_2(b)}\right)^{N-j} \left(\frac{d_2(b)}{2}\right)^{-\frac{j+1}{2}} \Gamma\left(\frac{j+1}{2}\right) - \sum_{j=0}^N \binom{N}{j} \left(-\frac{|d_1(b)|}{2d_2(b)}\right)^{N-j} \left(\frac{d_2(b)}{2}\right)^{-\frac{j+1}{2}} \gamma\left(\frac{j+1}{2}, \frac{d_1(b)^2}{8d_2(b)}\right) \right).$$

Rearranging terms, using $\zeta(b) = d_1(b)/(2\sqrt{d_2(b)})$ and the fact that $d_2(b) > 0$,

$$(B.18) \quad \varphi(b) = d_0 d_2(b)^{-\frac{N+1}{2}} e^{\frac{\zeta(b)^2}{2}} \left(\sum_{j=1,3,5,\dots}^N \binom{N}{j} 2^{\frac{j+1}{2}} (-\zeta(b))^{N-j} \Gamma\left(\frac{j+1}{2}\right) - \sum_{j=0}^N \binom{N}{j} 2^{\frac{j+1}{2}} (-|\zeta(b)|)^{N-j} \gamma\left(\frac{j+1}{2}, \frac{\zeta(b)^2}{2}\right) \right).$$

Combining the expressions for $\varphi(b)$ in (B.15) for even N and (B.18) for odd N , and using the above definitions for d_0 and $d_2(b)$ and $\zeta(b)$ yields the expression for $\varphi(b)$ in the proposition. \square

B.2 Proof of Proposition 2

Proof. The proof omits the subscript N from s_N^+ for brevity. When $\mu^2 = 0$, $\varphi(b)$ is a Cauchy distribution with unique mode σ_{wu}/σ_w^2 such that $s^+ = 0$. When $\rho^2 = 0$, $\varphi(b)$ has single mode at zero. The remainder of the proof assumes $\mu^2, \rho^2 > 0$.

The derivative of the density function $\varphi(b)$ w.r.t b is

$$(B.19) \quad \frac{d\varphi(b)}{db} = \frac{\sigma_w^2 \varphi(b) (N+1)}{b' \Sigma_u^{-1} \sigma_{wu} - 1} \Sigma_u^{-1} \left[\frac{((b' \Sigma_u^{-1} b) \sigma_{wu} + (1 - b' \Sigma_u^{-1} \sigma_{wu}) b) g(\zeta(b))}{1 + (b - \sigma_{wu}/\sigma_w^2)' \Sigma_u^{-1} \sigma_w^2 (b - \sigma_{wu}/\sigma_w^2)} - \frac{\sigma_{wu}}{\sigma_w^2} \right],$$

where $\zeta(b)$ is defined in Proposition 1 and

$$(B.20) \quad g(t) = 1 + \frac{1}{N+1} \frac{dh(t)}{dt} \frac{t}{h} = \begin{cases} 1 + \frac{1}{N+1} (t^2 + t g_e(t)) & \text{if } N \text{ is even;} \\ 1 + \frac{1}{N+1} (t^2 + t g_o(t)) & \text{if } N \text{ is odd;} \end{cases}$$

where

$$g_e(t) = \frac{\sum_{j=0,2,4,\dots}^{N-2} \binom{N}{j} (N-j) (t/\sqrt{2})^{N-j-1} \Gamma\left(\frac{j+1}{2}\right) / \sqrt{2}}{\sum_{j=0,2,4,\dots}^N \binom{N}{j} (t/\sqrt{2})^{N-j} \Gamma\left(\frac{j+1}{2}\right)};$$

$$g_o(t) = \frac{\sum_{j=1,3,5,\dots}^{N-2} \binom{N}{j} (N-j) (t/\sqrt{2})^{N-j-1} \Gamma\left(\frac{j+1}{2}\right) / \sqrt{2} + \text{sgn}(t) N \sum_{j=0}^{N-1} \binom{N-1}{j} (-|t|/\sqrt{2})^{N-j-1} \gamma\left(\frac{j+1}{2}, \frac{t^2}{2}\right) / \sqrt{2}}{\sum_{j=1,3,5,\dots}^N \binom{N}{j} (t/\sqrt{2})^{N-j} \Gamma\left(\frac{j+1}{2}\right) - \sum_{j=0}^N \binom{N}{j} (-|t|/\sqrt{2})^{N-j} \gamma\left(\frac{j+1}{2}, \frac{t^2}{2}\right)}.$$

The function $g(t)$ depends only on N , and is plotted in Figure B.2.

At any mode b_m , the term in square brackets in (B.19) is equal to zero. This means that

$$(B.21) \quad \frac{(b'_m \Sigma_u^{-1} b_m) \sigma_{wu} + (1 - b'_m \Sigma_u^{-1} \sigma_{wu}) b_m}{1 + (b_m - \sigma_{wu}/\sigma_w^2)' \Sigma_u |w| / \sigma_w^2 (b_m - \sigma_{wu}/\sigma_w^2)} g(\zeta(b_m)) = \frac{\sigma_{wu}}{\sigma_w^2}.$$

Since $g(t)$ is a scalar, the $N \times 1$ vector in front of $g(\zeta(b_m))$ in (B.21) is proportional to σ_{wu} . That vector is a linear combination of b_m and σ_{wu} with scalar weights, which implies that b_m is proportional to σ_{wu} as well. Let

$$(B.22) \quad b_m = \frac{\sigma_{wu}/\sigma_w^2}{1 + s},$$

where $s \in \mathbb{R}$ is an unknown scalar determining the proportionality to σ_{wu} corresponding to a mode of $\varphi(b)$. The proof proceeds by finding s^+ , or the value of s that corresponds to the major mode $b^+ = \operatorname{argmax} \varphi(b)$.

Substituting (B.22) into (B.21),

$$(B.23) \quad \frac{(1 - \rho^2)(1 + s)}{s^2 + (1 - \rho^2)(1 + 2s)} g \left(-\frac{(1 - \rho^2 + s)|\mu|/\sqrt{1 - \rho^2}}{\sqrt{s^2 + (1 - \rho^2)(1 + 2s)}} \right) \frac{\sigma_{wu}}{\sigma_w^2} = \frac{\sigma_{wu}}{\sigma_w^2}.$$

Any value of s corresponding to a mode, including s^+ , must therefore be a root of the scalar equation

$$(B.24) \quad \frac{(1 - \rho^2)(1 + s)}{s^2 + (1 - \rho^2)(1 + 2s)} g \left(-\frac{(1 - \rho^2 + s)|\mu|/\sqrt{1 - \rho^2}}{\sqrt{s^2 + (1 - \rho^2)(1 + 2s)}} \right) = 1.$$


Let $s = (1 - \rho^2)(k - 1)$. With this transformation of variables, (B.24) can be rewritten as

$$(B.25) \quad \frac{(1 - \rho^2)k + \rho^2}{(1 - \rho^2)k^2 + \rho^2} g(r(k)) = 1, \quad \text{where } r(k) = -\frac{k|\mu|/\sqrt{1 - \rho^2}}{\sqrt{k^2 + \rho^2/(1 - \rho^2)}}.$$

When $\mu^2, \rho^2 > 0$, the function $r(k)$ is strictly decreasing and its range is $(-|\mu|/(1 - \rho^2)^{\frac{1}{2}}, |\mu|/(1 - \rho^2)^{\frac{1}{2}})$. Equation (B.25) can be rewritten as the quadratic equation

$$(B.26) \quad k^2 - g(r(k))k - \frac{\rho^2}{1 - \rho^2} (g(r(k)) - 1) = 0,$$

FIGURE B.1: Fixed Points of κ_1 and κ_2



Notes: Illustrative examples for $N = 2$, $\rho^2 = 0.9^2$

with solutions

$$(B.27) \quad \begin{aligned} \kappa_1(k) &= \frac{1}{2}g(r(k)) + \frac{1}{2}\sqrt{g(r(k))^2 + \frac{4\rho^2}{1-\rho^2}(g(r(k)) - 1)}; \\ \kappa_2(k) &= \frac{1}{2}g(r(k)) - \frac{1}{2}\sqrt{g(r(k))^2 + \frac{4\rho^2}{1-\rho^2}(g(r(k)) - 1)}. \end{aligned}$$


Values of k that solve (B.25) are therefore fixed points of either one of the following two recurrence relations:

$$(B.28) \quad k_{n+1} = \kappa_1(k_n), \quad k_{n+1} = \kappa_2(k_n).$$

Figure B.1 provides illustrative examples of both relationships.

Consider $k_{n+1} = \kappa_1(k_n)$. It is easy to verify that $\kappa_1(0) = 1$ since $r(0) = 0$ and $g(0) = 1$. In addition, $\kappa_1(k) > 1$ for $k \neq 0$, since $g(t) > 1$ for $t \neq 0$ and $\rho^2/(1-\rho^2) > 0$ (see left panel of Figure B.2). Since $\kappa_1(k) \geq 1$, $k_{n+1} = \kappa_1(k_n)$ has no fixed points on the interval $(-\infty, 0]$. Note that $\text{sgn}(d\kappa_1(k)/dk) = \text{sgn}(dg(r(k))/dr(k)) \cdot \text{sgn}(dr(k)/dk)$. It is straightforward to check that $dg(t)/dt < 0$ for $t < 0$. Since $r(k) < 0$ for $k > 0$ and $dr(k)/dk < 0$, it follows that $d\kappa_1(k)/dk > 0$ for $k > 0$.

FIGURE B.2: The $g(t)$, $\varphi_k(k)$ and $f_k(k)$ functions.



Next, note that $\lim_{k \rightarrow \infty} \kappa_1(k) = \bar{k}$, where

$$(B.29) \quad \bar{k} = \frac{1}{2}g\left(\frac{|\mu|}{(1-\rho^2)^{\frac{1}{2}}}\right) + \frac{1}{2}\sqrt{g\left(\frac{|\mu|}{(1-\rho^2)^{\frac{1}{2}}}\right)^2 + \frac{4\rho^2}{1-\rho^2}\left(g\left(\frac{|\mu|}{(1-\rho^2)^{\frac{1}{2}}}\right) - 1\right)}.$$

Since $\kappa_1(k)$ is continuous, strictly increasing for $k > 0$, and \bar{k} is a finite positive constant, there must be some value $0 < k' < \bar{k}$ such that $\kappa_1(k') - k' < 0$. Because $\kappa_1(0) - 0 > 0$, by the intermediate value theorem, there exists at least one fixed point of $k_{n+1} = \kappa_1(k_n)$ within the interval $[0, k']$. Since $k = 0$ is not a fixed point as $\kappa_1(0) = 1 > 0$, any fixed point must be strictly larger than one and strictly smaller than \bar{k} .

Let $k^+ > 0$ be the largest fixed point, $k^+ = \kappa_1(k^+)$. By definition, this means that $\kappa_1(k_n) < k_n$ for all $k_n \in (k^+, \infty)$, else there would be a fixed point that is strictly larger than k^+ . Therefore $\kappa_1(k_n) - \kappa_1(k^+) < k_n - k^+$ is a contraction and $\lim_{n \rightarrow \infty} k_n = k^+$ for any starting point $k_0 \in [k^+, \infty)$. The limit in (B.29) above is guaranteed to lie in (k^+, ∞) , and therefore the starting point $k_0 = \bar{k}$ guarantees convergence to k^+ . Finally, since $\kappa_1(k^+) > 1$, it must be the case that $s^+ = (1 - \rho^2)(k^+ - 1) \in (0, \bar{k})$.

Consider $k_{n+1} = \kappa_2(k_n)$. It is easy to verify that $\kappa_2(0) = 0$ since $r(0) = 0$ and $g(0) = 1$. Therefore, $k = 0$ is a fixed point of $k_{n+1} = \kappa_2(k_n)$. In addition,

$\kappa_2(k) < 0$ for $k \neq 0$, since $g(t) > 1$ for $t \neq 0$ and $\rho^2/(1 - \rho^2) > 0$. Since $\kappa_2(k) < 0$, $k_{n+1} = \kappa_2(k_n)$ has no fixed points within the interval $(0, \infty)$, and therefore $k_{n+1} = \kappa_2(k_n)$ has no fixed points larger than k^+ .

It remains to be shown that $s^+ = (1 - \rho^2)(k^+ - 1)$ selects the major mode, and not any other mode of $\varphi(b)$. To see this, first substitute (B.22) into the expression for $\varphi(b)$ in Proposition 1 and use the previous transformation of variables to obtain the scalar function

$$(B.30) \quad \varphi_k(k) \propto \left(\frac{((1 - \rho^2)k + \rho^2)^2}{(1 - \rho^2)k^2 + \rho^2} \right)^{\frac{N+1}{2}} h(r(k)).$$

Intuitively, $\varphi_k(k)$ provides the likelihood of values of b that are proportional to σ_{wu} , where these values are indexed by $k = 1 + s/(1 - \rho^2)$, the scalar that determines the constant of proportionality. As shown above, these values must include all modes of $\varphi(b)$.

First, suppose $k^- < 0$ corresponds to a mode. Since $k^- < -k^-$,

$$(B.31) \quad \begin{aligned} (1 - \rho^2)k^- + \rho^2 &< (1 - \rho^2)(-k^-) + \rho^2 \\ \Rightarrow \frac{((1 - \rho^2)k^- + \rho^2)^2}{(1 - \rho^2)(k^-)^2 + \rho^2} &< \frac{((1 - \rho^2)(-k^-) + \rho^2)^2}{(1 - \rho^2)(-k^-)^2 + \rho^2}. \end{aligned}$$

In addition, it is easy to check that $h(r(k^-)) = h(r(-k^-))$. Therefore $f_k(-k^-) > f_k(k^-)$, such that no $k^- < 0$ can be the major mode.

Next, we show that no $0 \leq k < k^+$ can be the major mode. Since at any mode $\frac{k + \rho^2/(1 - \rho^2)}{k^2 + \rho^2/(1 - \rho^2)} g(r(k)) = 1$, the function

$$(B.32) \quad f_k(k) = ((1 - \rho^2)k + \rho^2)^{\frac{N+1}{2}} h(r(k)) g(r(k))^{-\frac{N+1}{2}}$$

intersects $\varphi_k(k)$ at all modes, as illustrated in the right panel of Figure B.2. It is straightforward to check that $h(r(k))/g(r(k))$ is strictly increasing for $k \in [0, \infty)$, and therefore $f_k(k)$ is also strictly increasing for $k \in [0, \infty)$. Since k^+ is the largest value for which $f_k(k) = \varphi_k(k)$, k^+ must be the major mode.

Next, we establish how s^+ varies with ρ^2 , μ^2 and N . Define $L(k, \rho^2, \mu^2) =$

$k - \kappa_1(k)$. Using $L(k^+, \rho^2, \mu^2) = 0$ and the implicit function theorem,

$$(B.33) \quad \frac{dk^+}{d\rho^2} = -\frac{\partial L}{\partial \rho^2}(k^+, \rho^2, \mu^2) / \frac{\partial L}{\partial k}(k^+, \rho^2, \mu^2)$$

where

$$\begin{aligned} \frac{\partial L}{\partial \rho^2}(k^+, \rho^2, \mu^2) &= -\frac{k^+ - 1}{1 - \rho^2} \cdot \frac{1 + k^+ - g(t^+) + \frac{1}{2}(1 + k^+)\epsilon_g(t^+)}{2k^+ - g(t^+)}, \\ \frac{\partial L}{\partial k}(k^+, \rho^2, \mu^2) &= 1 - \frac{\rho^2}{2k^+ - g(t^+)} \frac{\epsilon_g(t^+)}{(1 - \rho^2)k^+}, \end{aligned}$$

and $t^+ = r(k^+)$ and $\epsilon_g(t) = \frac{dg}{dt}(t)/(g(t)/t)$ is the elasticity of $g(t)$. Therefore,

$$(B.34) \quad \begin{aligned} \frac{ds^+}{d\rho^2} &= (1 - \rho^2) \frac{dk^+}{d\rho^2} - \frac{k^+ - 1}{1 - \rho^2} \\ &= (k^+ - 1) \left(\frac{k^+(k^+ + 1)}{2} + \frac{\rho^2}{1 - \rho^2} \right) \\ &\quad \cdot \left(\frac{\epsilon_g(t^+) - 2\frac{g(t^+)-1}{g(t^+)+1}}{g(t^+)k^+ + \frac{\rho^2}{1-\rho^2} (2(g(t^+) - 1) - \epsilon_g(t^+))} \right), \end{aligned}$$

where the last expression follows from substituting in (B.25) and rearranging. Note that $k^+ - 1 > 1$, and it is easy to verify that $\epsilon_g(0) = 0$ and $2(g(t) - 1)/(g(t) + 1) < \epsilon_g(t) < 2(g(t) - 1)$ for all $t \neq 0$, as illustrated in the left and right panels of Figure B.3. Therefore, $ds^+/d\rho^2 > 0$ and s^+ is strictly increasing in ρ^2 . Similarly, using the implicit function theorem,

$$(B.35) \quad \frac{dk^+}{d\mu^2} = -\frac{\partial L}{\partial \mu^2}(k^+, \rho^2, \mu^2) / \frac{\partial L}{\partial k}(k^+, \rho^2, \mu^2),$$


where

$$\frac{\partial L}{\partial \mu^2}(k^+, \rho^2, \mu^2) = -\frac{k^+(k^+ - 1)}{2k^+ - g(t^+)} \cdot \frac{g(t^+)\epsilon_g(t^+)/\mu^2}{2(g(t^+) - 1)},$$

such that

$$(B.36) \quad \frac{ds^+}{d\mu^2} = (1 - \rho^2) \frac{dk^+}{d\mu^2},$$

FIGURE B.3: Bounds on $\epsilon_g(t)$.



$$= \frac{(1 - \rho^2)g(t^+)(k^+)^2(k^+ - 1)\epsilon_g(t^+)/\mu^2}{2(g(t^+) - 1) \left((2k^+ - g(t^+))k^+ - \frac{\rho^2}{1 - \rho^2}\epsilon_g(t^+) \right)}.$$

Since $k^+ - 1 > 0$ and $g(t^+) - 1 > 0$, and

$$(B.37) \quad 2(k^+ - g(t^+))k^+ > \frac{\rho^2}{1 - \rho^2}\epsilon_g(t^+) \Leftrightarrow 2(g(t^+) - 1) > \epsilon_g(t^+),$$

it follows that $ds^+/d\mu^2 > 0$ and s^+ is strictly increasing in μ^2 . Finally, (B.27) implies that k^+ is increasing in $g(t)$ and since $g(t)$ is decreasing in N (left panel of Figure B.2), it follows that s^+ is strictly decreasing in N for $\mu^2 > 0$. \square

B.3 Proof of Proposition 3

Proof. If Θ is full row rank but not full column rank, i.e. $\text{rank}(\Theta) = H < N$, then the distribution of $[\nu_1, (\Theta\nu_2)']$ is jointly normal as in (11), but with (positive definite) covariance matrix $\tilde{\Sigma}_{wu}$. In this case, Propositions 1 and 2 apply directly with the substitutions given in the Proposition.

If Θ has full column rank, $d = \Theta b$ is a change of variables that can be exactly inverted using the Moore-Penrose pseudoinverse despite Θ not being square, $b = \Theta^\dagger d$, where $\Theta^\dagger = (\Theta'\Theta)^{-1}\Theta'$. $\Theta : \mathbb{R}^K \rightarrow \mathbb{R}^H$, $H > N$, is a linear

map, and applying Theorem 4.7 of Negro (2024) for the pdf induced for d yields

$$(B.38) \quad \varphi_\delta(d) = \det(\Theta' \Theta)^{-1/2} \varphi(\Theta^\dagger d),$$

where $\varphi_\delta(d)$ is defined with respect to the Hausdorff measure \mathcal{H}^N for \mathcal{H}^N -almost everywhere $d \in \Theta \mathbb{R}^N$. \square

B.4 Proof of Proposition 4

Proof. Given Assumptions 1-4, the OLS-based IRF estimators, $\hat{\delta}_{OLS}^{LP}$ and $\hat{\delta}_{OLS}^{VAR}$, are consistent. With identical estimands, both also have a common asymptotic OLS bias $d_{OLS} = \Theta \sigma_{wu}^{VAR} / (\sigma_w^{VAR})^2 = \sigma_{wu}^{LP} / (\sigma_w^{LP})^2$. By absolute homogeneity of norms, Proposition 3 implies that the difference in modal IV bias depends only on the relative magnitude of s_R^+ . If the LP and VAR models share the same first-stage equation, then $\mu_{VAR}^2 = \mu_{LP}^2 = \mu^2$. So, s_R^+ only differs because of differences in R or ρ^2 . It is always true that $\rho_{LP}^2 \leq \rho_{VAR}^2$, since

$$(B.39) \quad \begin{aligned} \rho_{LP}^2 &= (\sigma_{wu}^{LP})' (\Sigma_u^{LP})^{-1} \sigma_{wu}^{LP} / (\sigma_w^{LP})^2, \\ &= (\sigma_{wu}^{VAR})' \Theta' (\Sigma_u^{LP})^{-1} \Theta \sigma_{wu}^{VAR} / (\sigma_w^{VAR})^2, \\ &= (\sigma_{wu}^{VAR})' \Theta' (\Theta \Sigma_u^{VAR} \Theta' + \Omega)^{-1} \Theta \sigma_{wu}^{VAR} / (\sigma_w^{VAR})^2, \\ &\leq (\sigma_{wu}^{VAR})' \Theta' (\Theta \Sigma_u^{VAR} \Theta')^\dagger \Theta \sigma_{wu}^{VAR} / (\sigma_w^{VAR})^2 = \rho_{VAR}^2 \end{aligned}$$

and $\Omega = \Omega_1 (\mathcal{I}_{H-1} \otimes \Sigma_{wv}^{VAR}) \Omega_1'$ is a positive semidefinite matrix.¹ Part (i) follows since $R^{VAR} \leq R^{LP}$ and $\rho_{VAR}^2 \geq \rho_{LP}^2$, and the constant of proportionality in the mode is strictly increasing in R and strictly decreasing in ρ^2 (see Proposition 2). For strictly positive values of $\rho_{VAR}^2 \geq \rho_{LP}^2$ and μ^2 , if $H \geq K$ then $R^{VAR} < R^{LP}$, implying part (ii). Since $R^{VAR} \leq R^{LP}$, $\rho_{VAR}^2 > \rho_{LP}^2$ implies part (iii) for any values of H and K . \square

¹For example, when $\hat{\delta}_{IV}^{LP}$ and $\hat{\delta}_{IV}^{VAR}$ contain impulse responses over horizons $0, \dots, H-1$, then

$$\Omega_1 = \begin{bmatrix} 0 & 0 & \dots & 0 \\ [\theta \ \Theta]_1 & 0 & \dots & 0 \\ [\theta \ \Theta]_2 & [\theta \ \Theta]_1 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ [\theta \ \Theta]_{H-1} & [\theta \ \Theta]_{H-2} & \dots & [\theta \ \Theta]_1 \end{bmatrix}, \text{ where } [\theta \ \Theta]_i \text{ is the } i\text{-th row of } [\theta \ \Theta].$$

B.5 Proof of Proposition 5

The highest order polynomial term in the function $g(t)$ is $\frac{1}{N+1}t^2$, see (B.20). As a result, $\lim_{t \rightarrow \infty} g(t)/t^2 = 1/(N+1)$, and therefore

$$(B.40) \quad \begin{aligned} \lim_{\rho^2 \rightarrow 1} (1 - \rho^2) g \left(-\frac{(1 - \rho^2 + s)|\mu|/\sqrt{1 - \rho^2}}{\sqrt{s^2 + (1 - \rho^2)(1 + 2s)}} \right) \\ = \lim_{\rho^2 \rightarrow 1} (1 - \rho^2) g \left(-\frac{|\mu|}{\sqrt{1 - \rho^2}} \right) = \frac{\mu^2}{N+1} \end{aligned}$$

Taking limits of (B.24) and setting $N = R$,

$$(B.41) \quad \left(1 + \lim_{\rho^2 \rightarrow 1} s\right) / \left(\lim_{\rho^2 \rightarrow 1} s\right)^2 \cdot \frac{\mu^2}{R+1} = 1.$$

Solving for the positive root corresponding to the major mode,

$$(B.42) \quad \lim_{\rho^2 \rightarrow 1} s^+ = \frac{1}{2} \frac{\mu^2}{R+1} + \frac{1}{2} \left(\left(\frac{\mu^2}{R+1} + 2 \right)^2 - 4 \right)^{\frac{1}{2}}.$$

Substituting into $\lim_{\rho^2 \rightarrow 1} B(\rho^2, \mu^2) = 1/(1 + \lim_{\rho^2 \rightarrow 1} s^+)$ yields

$$(B.43) \quad \lim_{\rho^2 \rightarrow 1} B(\rho^2, \mu^2) = \left(1 + \frac{1}{2} \frac{\mu^2}{R+1} + \frac{1}{2} \left(\left(\frac{\mu^2}{R+1} + 2 \right)^2 - 4 \right)^{\frac{1}{2}} \right)^{-1} \leq 1$$

Setting $\lim_{\rho^2 \rightarrow 1} B(\rho^2, m(\tau)) = \tau$ and solving for $m(\tau)$ yields $m(\tau) = (R+1) \frac{(1-\tau)^2}{\tau}$.

We know that $\sup_{\rho^2} \{B(\rho^2, \mu^2)\} \geq \lim_{\rho^2 \rightarrow 1} B(\rho^2, \mu^2)$, with equality when $\mu^2 \geq \underline{\mu}^2$.

Since $B(\rho^2, \mu^2) = |\rho|/(1 + s_R^+)$ is strictly decreasing in μ^2 (Proposition 2), all models with $\mu^2 \leq m(\tau)$ must have $B(\rho^2, \mu^2) \geq \tau$ and, if $\mu^2 \geq \underline{\mu}^2$, all models with $\mu^2 > m(\tau)$ must have $B(\rho^2, \mu^2) < \tau$.

B.6 Proof of Proposition 6

Proof. The result follows trivially from (16) since $C(m(\tau), \alpha) \geq C(\mu^2, \alpha)$ for all $\mu^2 \leq m(\tau)$. \square

B.7 Proof of Proposition 7

The existence and uniqueness of a continuous differentiable function $m^c(\tau)$ follow from the Implicit Function Theorem and the fact that the derivative of $B = |\bar{\rho}/(1 + s_R^+)|$ is strictly negative and bounded away from zero (Proposition 2). The remainder of the proof is analogous to that of Proposition 5.

B.8 Proof of Proposition 8

Proof. Since $\hat{\Sigma}_{wv} \xrightarrow{p} \Sigma_{wv}$ (Assumption 4), by the Continuous Mapping Theorem and (13), $\hat{\rho}^2 \rightarrow \bar{\rho}^2$, and from Proposition 7,

$$(B.44) \quad \hat{m}^c(\tau) \xrightarrow{p} m^c(\tau), \quad \min\{\hat{m}^c(\tau), m^c(\tau)\} \xrightarrow{p} m^c(\tau).$$

As in Proposition B.6, we have from $F \xrightarrow{d} \chi_1^2(m^c(\tau))$ that,

$$(B.45) \quad \sup_{\mu^2 \leq m^c(\tau)} \lim_{T \rightarrow \infty} \text{Prob}(F > C(m^c(\tau), \alpha)) = \alpha.$$

Now, consider

$$(B.46) \quad \text{Prob}(F > C(\hat{m}^c(\tau), \alpha)) \leq \text{Prob}(F > \min(C(\hat{m}^c(\tau), \alpha), m^c(\tau))), \\ \xrightarrow{p} \text{Prob}(F > C(m^c(\tau), \alpha)) = \alpha.$$

Since the critical value function $C(\mu^2, \alpha)$ is weakly increasing in μ^2 , we have

$$(B.47) \quad C(m^c(\tau), \alpha) \geq C(\mu^2, \alpha) \quad \text{for all } \mu^2 \leq m^c(\tau).$$

Thus,

$$(B.48) \quad \lim_{T \rightarrow \infty} \text{Prob}(F > C(m^c(\tau), \alpha)) \leq \alpha \quad \text{for all } \mu^2 \leq m^c(\tau),$$

and the statement in the Proposition follows. □

C Bootstrap Implementation and Simulation Evidence

This section outlines a parametric bootstrap implementation of the weak instrument test, and presents simulation evidence on the performance of the baseline weak instrument test described in Section 4.2 in terms of size and power. The procedure for obtaining the bootstrap critical value $C^{bs}(\alpha, \tau)$ at significance level α and tolerance level τ is as follows:

1. Regress the scaling (first-stage) variable x_t^Y on X_{t-1} (a constant and lagged controls) using OLS. Obtain the fitted values \hat{x}_t^Y and the residual variance estimate $\hat{\sigma}_Y^2$.
2. Define $\mu^2 = (R+1)(1-\tau)^2/\tau$, where R is the rank of the impulse response (see Section 2.3). Compute

$$\sigma_w^2 = \frac{\hat{\sigma}_Y^2}{1 + \mu^2/T}, \quad \Pi = \sigma_w \sqrt{\mu^2/T}.$$

3. For each bootstrap iteration $b = 1, \dots, N_b$:
 - (a) Draw residuals $w_t^b \sim \mathcal{N}(0, \sigma_w^2)$ for $t = 1, \dots, T$, and draw instruments $\tilde{Z}_t^b \sim \mathcal{N}(0, 1)$. Orthogonalize \tilde{Z}_t^b with respect to X_{t-1} , standardize to have unit sample covariance, and denote the result by Z_t^b . Construct

$$x_t^{Y,b} = \Pi Z_t^b + \hat{x}_t^Y + w_t^b, \quad t = 1, \dots, T.$$

- (b) Regress $x_t^{Y,b}$ on Z_t^b on X_{t-1} . Record the F -statistic, F^b , associated with the OLS estimate of Π .
4. Compute the bootstrap critical value $C^{bs}(\alpha, \tau)$ as the upper α -quantile of the empirical distribution of $\{F^b\}_{b=1}^{N_b}$.

Table C.1 reports empirical rejection rates for nominal 5% tests of the null hypothesis of a weak instrument with bias tolerance $\tau = 0.10$ for a range of DGPs. In each case, the test is based on the joint bias in the impulse response over multiple horizons. We consider impulse responses from several DGPs including an impulse response in a simple AR(1) model with persistence 0.90 and the other parameters such that $\rho^2 = 0.90$; the inflation response to an interest rate shock in a textbook New Keynesian model as in Section 1.1 and Appendix A; and responses from several empirical DGPs calibrated to actual

TABLE C.1: EMPIRICAL SIZE OF NOMINAL 5% TESTS FOR VARIOUS DGPs

	SVAR-IV			LP-IV		
	T small		T large	T small		T large
	bs	asy	asy	bs	asy	asy
AR(1) Model	5.4	6.0	4.9	5.5	7.9	5.3
New Keynesian Model	5.3	5.7	5.1	5.3	6.8	4.9
MP: Gertler and Karadi (2015)	4.7	2.4	4.9	5.5	1.5	5.1
OIL: Montiel Olea et al. (2021)	4.8	2.0	5.1	4.9	0.5	4.5
UNC: Carriero et al. (2015)	5.3	1.4	4.5	5.0	0.2	4.5
TAX: Mertens and Montiel Olea (2018)	4.7	2.8	4.6	5.4	3.1	4.8

Note: Rejection rates of the joint weak instrument test with $\tau = 0.10$ and $\alpha = 0.05$ across 10,000 Monte Carlo samples under bootstrapped (bs; 5,000 samples) and asymptotic (asy) critical values for the F -statistic. For the small samples, $T = 250$ for the AR(1) and NK models. For the other DGPs, T equals the empirical sample sizes in the papers cited. For the large sample simulations, $T = 10,000$.

applications with U.S. time series in the literature. Specifically, we consider the response of industrial production to a monetary policy shock as in Gertler and Karadi (2015); the response of industrial production to an uncertainty shock as in Carriero et al. (2015); the response of global economic activity to an oil supply shock as in Montiel Olea et al. (2021); and the response of real GDP to a marginal tax rate shock as in Mertens and Montiel Olea (2018). For the AR(1) and New Keynesian models, we set the sample size to $T = 250$ and the horizon to $H = 12$ and $H = 8$ periods, respectively. The monetary policy, uncertainty and oil shock applications use monthly samples of the same size as in the actual data (around 400-500 months), and we consider a four-year horizon, $H = 48$. The tax shock application uses an annual postwar sample of 65 years, and we consider a response horizon of $H = 6$ years. These values for T and H are representative for typical macro applications with U.S. quarterly, monthly, or annual data, and the corresponding results are labeled under the header “ T small”. We also consider results in simulations for much larger sample sizes (“ T large”) to verify the asymptotic validity of our testing procedures.

The rejection rates in Table C.1 are for 10,000 Monte Carlo samples. In all models, the concentration parameter μ^2 is set such that the true (joint) worst-case bias in the DGP is exactly equal to the tolerance level $\tau = 0.10$. For small

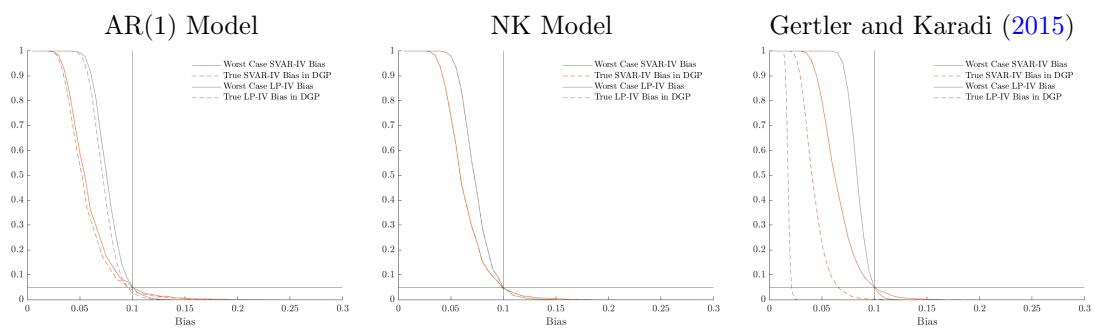
T , the table reports rejection rates based on bootstrapped and asymptotic critical values, the latter coming from the non-central χ_1^2 distribution. For T large, only those for the asymptotic critical values are shown. Finally, the tests are conducted with both the SVAR-IV and LP-IV impulse response estimators.

Consistent with the theory, Table C.1 shows that the empirical rejection rates in large samples are very close to the 5% nominal size of the test. In small samples, some size distortions inevitably emerge because the (asymptotic) approximation of the first-stage and reduced form OLS estimates as jointly Gaussian becomes less accurate, and because of the greater influence of estimation error in the other coefficients of the models (both the VARs and LPs include lags of various endogenous variables as additional regressors). In realistically sized samples, the performance of the test under the asymptotic χ_1^2 critical values nevertheless remains reasonably good.

The first and fourth columns in Table C.1 show that using the bootstrap procedure described in Appendix C instead of the asymptotic critical value results in meaningful improvements in performance, with rejection rates that are much closer to the nominal size overall. Based on these improvements in samples that are more typical for actual empirical applications, we recommend using the bootstrap to obtain the critical values in practice.

Figure C.4 evaluates the power of the weak instrument test for $\alpha = 0.05$ and $\tau = 0.10$ across three DGPs. Each panel plots empirical rejection rates as a function of the bias criterion for $\rho^2 \rightarrow 1$, across a grid for the concentration parameter μ^2 (full lines). For reference, the figure panels also plot rejection rates as a function of the actual modal bias based on the true values of the structural parameters in the DGP (dashed lines). Both are shown for SVAR-IV (in red) and LP-IV (in black). Each gridpoint reflects rejection rates in 10,000 Monte Carlo samples based on bootstrapped critical values of the F -statistic across 5,000 replications.

FIGURE C.4: Power Curves In Simulations



Notes: Horizontal lines marks $\alpha = 0.05$, vertical lines mark $\tau = 0.10$. Rejection rates across a grid for μ^2 . Each gridpoint shows the rejection rate across 10,000 Monte Carlo samples based on bootstrapped critical values of the F -statistic, with 5,000 bootstrap samples for each Monte Carlo sample.