

Shocks or Shifts? Identifying Macroeconomic Impulse Responses with Instruments

Daniele Colombo*

IAAE 2026

Lisbon

June, 2026

Preview of the Talk

- We often think of IRFs as freely **rescalable**. This is broadly true for point estimates, but not for inference: different normalizations define different estimands.
- In SVAR-IV and LP-IV, instrument validity is usually stated in terms of **latent structural shocks**, while the estimators used are equivalent to IV estimators for shifts in observed variables.
- Indeed, for **impact-normalized IRFs**, the analysis can be formulated under weaker reduced-form IV assumptions, without imposing structural-shock restrictions such as invertibility.
- **One-standard-deviation IRFs** use the structural-shock restrictions more directly, delivering power gains and greater robustness to weak instruments.
- I propose a **shock-based relevance diagnostic** for one-standard-deviation IRFs.

SVAR:

$$\mathbf{A}(L)\mathbf{y}_t = \mathbf{u}_t$$

reduced-form VAR

$$\mathbf{u}_t = \Theta_0 \mathbf{w}_t$$

innovations spanned by shocks

$$\mathbf{y}_t = \Theta(L)\mathbf{w}_t$$

SVMA representation

$$\mathbb{E}[\mathbf{w}_t \mathbf{w}_t'] = \text{diag}(\sigma_{w_1}^2, \dots, \sigma_{w_K}^2)$$

orthogonality

Instrument:

$$\mathbb{E}[\mathbf{z}_t w_{1,t}] = \boldsymbol{\alpha} \neq \mathbf{0}$$

relevance

$$\mathbb{E}[\mathbf{z}_t \mathbf{w}_{2:K,t}'] = \mathbf{0}$$

exogeneity

\implies IRFs to $w_{1,t}$ $\Theta(L)\mathbf{e}_1 = \mathbf{A}(L)^{-1}\Theta_0\mathbf{e}_1$ identified up to scale.

Standard approach: recover the impact vector using the SVAR-IV estimator of [Stock and Watson \(2018\)](#).

Motivating Example

Colombo and Toni (2025) construct an external instrument for a natural gas supply shock.

If we instrument gas **price**, we obtain:




Figure: Responses of gas prices and quantities. Other variables omitted. Robust $F = 12.92$.

Motivating Example

Colombo and Toni (2025) construct an external instrument for a natural gas supply shock.

If we instrument gas **quantity**, we obtain:

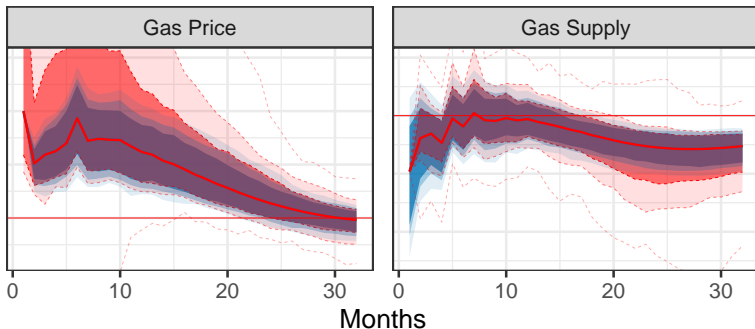


Figure: Responses of gas prices and quantities. Other variables omitted. Robust $F = 1.32$.

Different Normalizations Define Different Estimands

Suppose the target shock is $w_{1,t}$, with impact vector θ . The identifying moment condition is

$$\begin{aligned}\mathbb{E}[z_t \mathbf{u}'_t] &= \mathbb{E}[z_t \mathbf{w}'_t] \Theta'_0 \\ &= (\mathbb{E}[z_t w_{1,t}] \quad \mathbb{E}[z_t \mathbf{w}'_{2:K,t}]) \Theta'_0 \\ &= (\boldsymbol{\alpha} \quad \mathbf{0}) \Theta'_0 = \boldsymbol{\alpha} \theta' .\end{aligned}$$

$\implies \theta$ identified only up to scale, since not separately identified from $\boldsymbol{\alpha}$.

For the rest of the talk, I focus on the leading case of one scalar instrument for one structural shock.

Impact normalization on variable j :

$$\theta^{\text{UI}(j)} := \frac{\tilde{\theta}}{e'_j \tilde{\theta}}, \quad e'_j \theta^{\text{UI}(j)} = 1,$$

where $\tilde{\theta}$ is any nonzero rescaling of θ .

Different Normalizations Define Different Estimands

Impact normalization on variable j :

$$\boldsymbol{\theta}^{\text{UI}(j)} := \frac{\tilde{\boldsymbol{\theta}}}{e_j' \tilde{\boldsymbol{\theta}}} = \frac{\mathbb{E}[z_t \mathbf{u}_t]}{\mathbb{E}[z_t u_{j,t}]}, \quad e_j' \boldsymbol{\theta}^{\text{UI}(j)} = 1.$$

The corresponding plug-in estimator is

$$\hat{\boldsymbol{\theta}}^{\text{UI}(j)} = \frac{\sum_{t=1}^T z_t \hat{\mathbf{u}}_t}{\sum_{t=1}^T z_t \hat{u}_{j,t}}.$$

- Different choices of j define different estimands:

$$\boldsymbol{\theta}^{\text{UI}(\text{gas price})} \neq \boldsymbol{\theta}^{\text{UI}(\text{gas quantity})}.$$

- Hence, IRFs and confidence bands cannot be rescaled ex post, unless the scaling factor is fixed and non-random.
- The normalization is part of the estimator:

bootstrap \Rightarrow renormalize inside each draw.

Macro Identification as Standard IV

$$\widehat{\theta}^{UI(j)} = \frac{\sum_{t=1}^T z_t \widehat{\mathbf{u}}_t}{\sum_{t=1}^T z_t \widehat{u}_{j,t}}.$$

This is a ratio estimator:

$$\mathbb{E}[z_t u_{j,t}] \approx 0 \quad \Rightarrow \quad \text{weak-IV problem.}$$

Inference can be organized around the normalization condition, where weak-instrument concerns parallel weak-IV concerns (Montiel Olea et al., 2021).

- Delta-method confidence intervals only under strong-instrument asymptotics;
- first-stage F tests for instrument relevance based on $\mathbb{E}[z_t u_{j,t}] \neq 0$; robust versions allow for heteroskedasticity and serial dependence;
- weak-IV robust confidence sets based on Anderson–Rubin/grid inversion: the analytical MSW set (Montiel Olea et al., 2021) and the grid MBB AR set, where MBB percentile intervals robust to heteroskedasticity and serial dependence (Jentsch & Lunsford, 2022);

But... the Same Estimand/Estimator Arise Without Structural Shocks

Consider again

$$\mathbf{A}(L)\mathbf{y}_t = \mathbf{u}_t.$$

Interested in contemporaneous effect of $y_{j,t}$ on all variables in the system, conditional on \mathbf{X}_{t-1}

$$y_{i,t} = \beta_{ij}y_{j,t} + \boldsymbol{\Gamma}'_{ij}\mathbf{X}_{t-1} + v_{i,t}. \quad \forall i \neq j$$

By Frisch–Waugh–Lovell:

$$u_{i,t} = \beta_{ij}u_{j,t} + v_{i,t}, \quad \mathbb{E}[u_{j,t}v_{i,t}] \neq 0.$$

Let z_t an instrument satisfying:

$$\mathbb{E}[z_t u_{j,t}] = \pi \neq 0, \quad \mathbb{E}[z_t v_{i,t}] = 0.$$

Then

$$\mathbb{E}[z_t (u_{i,t} - \beta_{ij}u_{j,t})] = 0 \implies \beta_{ij} = \frac{\mathbb{E}[z_t u_{i,t}]}{\mathbb{E}[z_t u_{j,t}]}.$$

Therefore, the full unit-impact vector is

$$\begin{aligned} \mathbf{b}^{\text{UI}(j)} &= \frac{\mathbb{E}[z_t \mathbf{u}_t]}{\mathbb{E}[z_t u_{j,t}]}, & \hat{\mathbf{b}}^{\text{UI}(j)} &= \frac{\sum_{t=1}^T z_t \hat{\mathbf{u}}_t}{\sum_{t=1}^T z_t \hat{u}_{j,t}}. \\ \boldsymbol{\theta}^{\text{UI}(j)} &= \mathbf{b}^{\text{UI}(j)}, & \hat{\boldsymbol{\theta}}^{\text{UI}(j)} &= \hat{\mathbf{b}}^{\text{UI}(j)}. \end{aligned}$$

Structural Assumptions Imply Reduced-Form IV

If the structural model is true, the residual $v_{i,t}$ contains only non-target shocks. Hence,

$$\mathbb{E}[z_t \mathbf{w}'_{2:K,t}] = \mathbf{0} \quad \implies \quad \mathbb{E}[z_t v_{i,t}] = 0.$$

\implies structural exogeneity implies reduced-form exogeneity.

If the structural model is true and structural exogeneity holds, then

$$\mathbb{E}[z_t u_{j,t}] = \Theta_{0,j1} \mathbb{E}[z_t w_{1,t}].$$

\implies structural relevance implies reduced-form relevance if $\Theta_{0,j1} \neq 0$.

Implication: a weak first stage for $u_{j,t}$ may reflect either a weak instrument or a poor normalization variable.

Conversely, reduced-form relevance for at least one valid normalization variable, implies structural relevance.

Does the Structural Model Add Anything?

Both models start from the same dynamics for the observables.

Both deliver the same estimand as a function of reduced-form moments

The instrument assumptions are not identical, but the structural assumptions are stronger: when the structural model is true, they imply the corresponding reduced-form IV assumptions.

LPs analogue

The main difference is interpretation:

- **Reduced-form IV view:** response to an exogenous unit movement in $y_{j,t}$.
- **Structural view:** unit-impact normalization of a latent structural-shock IRF.
- The structural view imposes more:
latent shocks + orthogonality + invertibility.

Can these stronger assumptions buy us more than interpretation?

Standard-Deviation Normalization

Consider again the moment condition $\mathbb{E}[z_t \mathbf{u}'_t] = \alpha \boldsymbol{\theta}'$. Using $\boldsymbol{\Sigma}_u = \boldsymbol{\Theta}_0 \boldsymbol{\Sigma}_w \boldsymbol{\Theta}'_0$, we can write $\mathbb{E}[z_t \mathbf{u}'_t] \boldsymbol{\Sigma}_u^{-1} \mathbb{E}[\mathbf{u}_t z_t] = \alpha^2 \boldsymbol{\theta}' \boldsymbol{\Sigma}_u^{-1} \boldsymbol{\theta} = \frac{\alpha^2}{\sigma_{w_1}^2}$. Hence, the impact of a one standard deviation shock is given by

$$\boldsymbol{\theta} \sigma_{w_1} = \text{sgn}(\alpha) \frac{\boldsymbol{\theta}}{(\boldsymbol{\theta}' \boldsymbol{\Sigma}_u^{-1} \boldsymbol{\theta})^{1/2}}.$$

Therefore,

$$\boldsymbol{\theta}^{UV} := \pm \frac{\tilde{\boldsymbol{\theta}}}{(\tilde{\boldsymbol{\theta}}' \boldsymbol{\Sigma}_u^{-1} \tilde{\boldsymbol{\theta}})^{1/2}}, \quad w_{1,t}^{UV} := \boldsymbol{\theta}^{UV'} \boldsymbol{\Sigma}_u^{-1} \mathbf{u}_t, \quad \text{Var}(w_{1,t}^{UV}) = 1.$$

and the plug-in estimator is

$$\hat{\boldsymbol{\theta}}^{UV} = \pm \frac{\sum_{t=1}^T z_t \hat{\mathbf{u}}_t}{\left[\left(\sum_{t=1}^T z_t \hat{\mathbf{u}}_t \right)' \hat{\boldsymbol{\Sigma}}_u^{-1} \left(\sum_{t=1}^T z_t \hat{\mathbf{u}}_t \right) \right]^{1/2}}.$$

Standard-Deviation Normalization

- **No anchor variable.** No need to select an anchor response, or to require that the (unknown) true impact on that variable is nonzero.
- **Less exposed to weak-IV problems.** In impact-normalized IRFs weakness can arise when $\mathbb{E}[z_t u_{j,t}]$ is small for a particular j . SD-normalization instead uses the full covariance between the instrument and the innovation vector.
 - \implies same instrument may be weak for impact-normalized IRFs but informative for SD-IRFs!
- **Natural scale for shock-based objects.** Delivers the natural object for variance decompositions, historical decompositions, and recovered shock series.

Example: What Is a Monetary Policy Shock?

Random policy-rate shift

Instrument: observe some of the coin flips used by the central banker to randomly move the policy rate beyond the systematic rule.

- Channel: **only through the interest rate**
- Natural object: **impact-normalized IRF**
- Example: response to a 25bp increase in the policy rate
- Rate moves by construction, **no impact confidence band for normalized variable**; uncertainty concerns the responses of the other variables.

Latent policy shock

Instrument: observe imperfect measure of shift in expectations, communication, or the perceived reaction function.

- Channel: **latent shift**, not only the interest rate
- Natural object: **one-SD structural-shock IRF**
- Same 25bp move may reflect different underlying shocks
- No variable is fixed by normalization on impact, **uncertainty applies to all variables**, including the policy rate.

A Shock-Based Weak-IV Diagnostic

The model's relevance condition is directly about the target shock, not a particular normalization equation:

$$\mathbb{E}[\mathbf{z}_t w_{1,t}] \neq \mathbf{0}.$$

Since

$$w_{1,t} = \text{Proj}(z_t \mid \mathbf{u}_t) = \boldsymbol{\theta}' \boldsymbol{\Sigma}_u^{-1} \mathbf{u}_t,$$

I propose the following shock-based diagnostic:

1. Estimate the reduced-form VAR and obtain $\hat{\mathbf{u}}_t$ and $\hat{\boldsymbol{\Sigma}}_u$.
2. Estimate the impact direction $\hat{\boldsymbol{\theta}}$, up to scale.
3. Recover the shock index:

$$\hat{w}_{1,t} = \hat{\boldsymbol{\theta}}' \hat{\boldsymbol{\Sigma}}_u^{-1} \hat{\mathbf{u}}_t.$$

4. Regress $\hat{w}_{1,t}$ on \mathbf{z}_t and report the associated first-stage F -statistic.

Note that $F_{\text{shock}} \geq F_{\text{anchor}}(j) \quad \forall j.$

Standard-Deviation Normalization

Colombo and Toni (2025) with SD normalization:




Figure: Responses of gas prices and quantities. Other variables omitted. Robust $F = 14.37$.

Standard-Deviation Normalization

Colombo and Toni (2025) with SD normalization:





Figure: Responses of gas prices and quantities. Other variables omitted. Robust $F = 14.37$.

Monte Carlo Design

- Same stable trivariate VAR in all DGPs:



- The instrument is generated as

$$z_t = \bar{m} + \psi_T \varepsilon_{1,t} + v_t, \quad v_t \sim N(0, 1).$$

- Instrument strengths are calibrated at sample size $T = 360$ (30 years of monthly data). Chosen so that the normalized-IRF first-stage diagnostic $F_{iv} > 10$ rejects in 100%, 50%, 10%, and 5% of simulations:

$$\psi_{360} \in \{1, 0.24, 0.14, 0.11\}.$$

- For the large-sample design, $T = 3600$, use local-to-zero scaling,

$$\psi_T = \psi_{360} \sqrt{360/T},$$

so that $\sqrt{T}\psi_T$ is held fixed across sample sizes.

Simulation Objects

500 simulations, 1000 bootstrap replications each

For each DGP, compare coverage and median interval length.


The four lines are:

1. MBB percentile, impact-normalized (bootstrap, non-robust);
2. grid-MBB Anderson–Rubin, impact-normalized (bootstrap, weak-IV robust);
3. MSW analytical, impact-normalized (analytical, weak-IV robust);
4. MBB percentile, one-standard-deviation (bootstrap, non-robust).


⇒ AR/MSW robust intervals can be unbounded when the normalization is weak (infinite length).

Impact coverage and length, 68%, $T = 3600$

A. Coverage of contemporaneous impact



B. Median length of contemporaneous-impact confidence interval




—x— MBB normalized —o— grid MBB AR normalized —△— MSW normalized —□— MBB one-SD


► Full IRF coverage and length

Impact coverage and length, 68%, $T = 360$

A. Coverage of contemporaneous impact




B. Median length of contemporaneous-impact confidence interval




—x— MBB normalized —o— grid MBB AR normalized —△— MSW normalized —□— MBB one-SD

Impact coverage and length, 95%, $T = 3600$

A. Coverage of contemporaneous impact



B. Median length of contemporaneous-impact confidence interval




—x— MBB normalized —o— grid MBB AR normalized —△— MSW normalized —□— MBB one-SD


► Full IRF coverage and length

Impact coverage and length, 95%, $T = 360$

A. Coverage of contemporaneous impact



B. Median length of contemporaneous-impact confidence interval



—x— MBB normalized —o— grid MBB AR normalized —v— MSW normalized —s— MBB one-SD

► Full IRF coverage and length

Rejection Rates of First-Stage Diagnostics

DGP	ψ_T	$F_{iv} > 10$	$F_{shock} > 10$	$F_{SW2012} > 10$
1	1.000	100.0%	100.0%	100.0%
2	0.076	51.2%	92.8%	23.6%
3	0.044	9.0%	41.8%	0.0%
4	0.035	4.4%	26.8%	0.0%

Table: Share of simulations in which the first-stage statistic exceeds 10, $T = 3600$.

Simulation Takeaways





- **Impact-normalized inference** is sensitive to the normalization equation: when the proxy is weak for the impact variable, AR/MSW intervals become wide or unbounded.
- **One-standard-deviation IRFs** deliver shorter MBB intervals and are more robust to weak instruments, because inference does not hinge on a weak impact-normalization equation.
- The usual anchor-based diagnostics can understate relevance for one-SD IRFs; the proposed **shock-based diagnostic** remains informative for the latent shock targeted by the instrument.

Conclusions

Do we believe that the shock exists as an economic object?

- **If not**, and the goal is simply to trace the effects of an exogenous movement in an observed variable, then the reduced-form IV interpretation may be enough. In that case, imposing a full structural-shock model may add unnecessary assumptions without changing the estimand.
- **If yes**, then the structural model should do more than provide language. Standard-deviation normalization follows naturally: it avoids choosing an anchor variable, recovers a shock with a well-defined scale, and delivers the objects needed for shock-based analysis.
- This distinction **matters for inference**. Normalization is part of the estimator, not an after-the-fact rescaling: different normalizations define different estimands and different weak-IV problems.
- I propose a **shock-based relevance diagnostic**, based on relevance for the latent shock rather than for a chosen normalization variable, which better tracks the primitive model assumptions.

References I

-  Stock, J. H., & Watson, M. W. (2018). Identification and estimation of dynamic causal effects in macroeconomics using external instruments. *The Economic Journal*, 128(610), 917–948.
-  Montiel Olea, J. L., Stock, J. H., & Watson, M. W. (2021). Inference in structural vector autoregressions identified with an external instrument. *Journal of Econometrics*, 225(1), 74–87.
-  Jentsch, C., & Lunsford, K. G. (2022). Asymptotically valid bootstrap inference for proxy svdrs. *Journal of Business & Economic Statistics*, 40(4), 1876–1891.
-  Colombo, D., & Toni, F. (2025). Gas prices and the macroeconomy. Available at SSRN 5237587.

Appendix - Local Projections

LP-OLS:

$$y_{t+h} = \beta_h x_t + \mathbf{\Gamma}'_h \mathbf{X}_{t-1} + v_{t+h}, \quad x_t \text{ instrumented by } z_t.$$

In the same spirit as the shock-based approach, one can instead project directly on the instrument:

$$y_{t+h} = c_h z_t + \mathbf{\Gamma}'_h \mathbf{X}_{t-1} + e_{t+h}.$$

- This identifies the dynamic response only up to scale: the units are those of the instrument, not of a normalized structural shock.
- There is no ratio estimator, so standard OLS confidence intervals are valid for c_h .
- Interpretation is limited to the shape and timing of responses:
e.g. whether the response is persistent, hump-shaped, or changes sign.

A joint diagnostic based on all impact innovations is closer to the shock-based SVAR diagnostic, since it asks whether the instrument is informative for any innovation. But it still does not, by itself, separate instrument strength from the shock's impact profile across variables.

LP-IV

LP:

$$\mathbf{y}_{t+h} = \boldsymbol{\theta}_h w_{1,t} + \mathbf{u}_{t+h}$$

horizon- h LP equation

$$\mathbf{y}_t = \boldsymbol{\Theta}(L)\mathbf{w}_t$$

SVMA representation

$$\mathbb{E}[\mathbf{w}_t \mathbf{w}_t'] = \text{diag}(\sigma_{w_1}^2, \dots, \sigma_{w_K}^2)$$

orthogonality

Instrument:

$$\mathbb{E}[z_t w_{1,t}] = \alpha \neq 0$$

relevance

$$\mathbb{E}[z_t w_{k,t+\tau}] = 0, \quad (k, \tau) \neq (1, 0), \quad \tau \leq h$$

lead-lag exogeneity

$$\implies \mathbb{E}[z_t \mathbf{u}_{t+h}] = \mathbf{0}.$$

Thus the instrument is valid for the horizon- h LP written in terms of the latent target shock. However, $w_{1,t}$ is unobserved, so regression cannot be estimated directly. Identification also requires that the shock be normalized wrt an observed impact variable, say $y_{j,t} = \mathbf{e}'_j \mathbf{y}_t$.

LP-IV: Impact Normalization

To obtain an estimable equation, express the target shock in units of an observed impact variable:

$$w_{1,t}^{(j)} = \gamma_j w_{1,t}, \quad y_{j,t} := \mathbf{e}'_j \mathbf{y}_t = w_{1,t}^{(j)} + r_{j,t} = \gamma_j w_{1,t} + r_{j,t}.$$

$r_{j,t}$ is the part of $y_{j,t}$ not due to the target shock: systematic movements, measurement error, and other shocks. Under the SVMA,

$$r_{j,t} = \sum_{(\ell,k) \neq (0,1)} \mathbf{e}'_j \boldsymbol{\Theta}_\ell \mathbf{e}_k w_{k,t-\ell}.$$

Since $w_{1,t} = (y_{j,t} - r_{j,t})/\gamma_j$, the horizon- h LP becomes

$$\mathbf{y}_{t+h} = \boldsymbol{\theta}_h^{(j)} y_{j,t} + \mathbf{u}_{t+h}^{IV,(j)}, \quad \boldsymbol{\theta}_h^{(j)} = \frac{\boldsymbol{\theta}_h}{\gamma_j}, \quad \mathbf{u}_{t+h}^{IV,(j)} = \mathbf{u}_{t+h} - \boldsymbol{\theta}_h^{(j)} r_{j,t}.$$

The shock-level restrictions imply $\mathbb{E}[z_t r_{j,t}] = 0$, hence

$$\mathbb{E}[z_t y_{j,t}] = \gamma_j \mathbb{E}[z_t w_{1,t}] \neq 0, \quad \mathbb{E}[z_t \mathbf{u}_{t+h}^{IV,(j)}] = \mathbf{0}.$$

Therefore,

$$\boldsymbol{\theta}_h^{(j)} = \frac{\mathbb{E}[z_t \mathbf{y}_{t+h}]}{\mathbb{E}[z_t y_{j,t}]}, \quad \widehat{\boldsymbol{\theta}}_h^{IV,(j)} = \frac{\sum_t z_t \mathbf{y}_{t+h}}{\sum_t z_t y_{j,t}}.$$

Again... the Same Estimand/Estimator Arise Without Structural Shocks

$$\mathbf{y}_{t+h}(y) = \mathbf{a}_h + \boldsymbol{\tau}_h^{(j)} \mathbf{y} + \mathbf{v}_{t+h}^{(j)}.$$

Assume that z_t shifts $y_{j,t}$, but affects \mathbf{y}_{t+h} only through that shift:

$$\mathbb{E}[z_t y_{j,t}] \neq 0, \quad \mathbb{E}[z_t \mathbf{v}_{t+h}^{(j)}] = \mathbf{0}, \quad \mathbf{y}_{t+h}(z, y) = \mathbf{y}_{t+h}(y).$$

Using consistency, $\mathbf{y}_{t+h} = \mathbf{y}_{t+h}(y_{j,t})$, and centering z_t ,

$$\mathbb{E}\left[z_t \left(\mathbf{y}_{t+h} - \boldsymbol{\tau}_h^{(j)} y_{j,t}\right)\right] = \mathbf{0} \quad \implies \quad \boldsymbol{\tau}_h^{(j)} = \frac{\mathbb{E}[z_t \mathbf{y}_{t+h}]}{\mathbb{E}[z_t y_{j,t}]}.$$

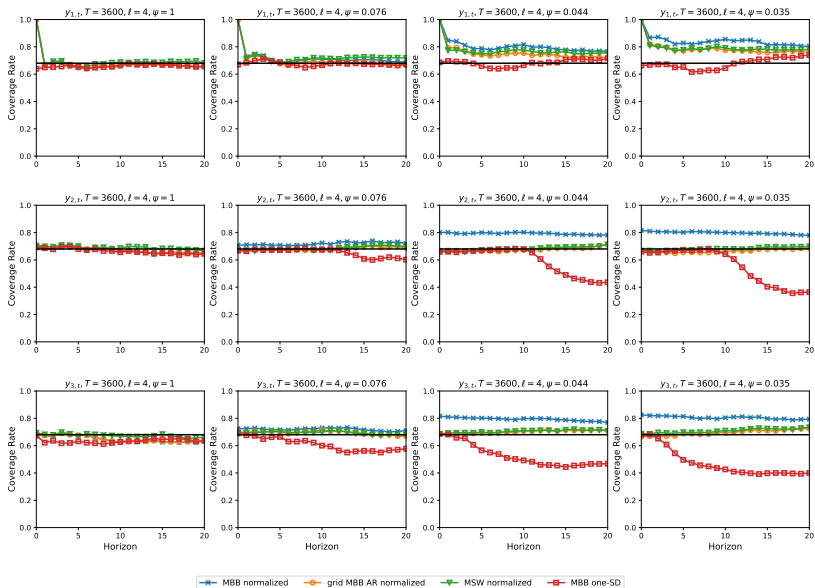
Therefore,

$$\boldsymbol{\tau}_h^{(j)} = \frac{\mathbb{E}[z_t \mathbf{y}_{t+h}]}{\mathbb{E}[z_t y_{j,t}]}, \quad \hat{\boldsymbol{\tau}}_h^{IV,(j)} = \frac{\sum_{t=1}^T z_t \mathbf{y}_{t+h}}{\sum_{t=1}^T z_t y_{j,t}}.$$
$$\implies \hat{\boldsymbol{\tau}}_h^{IV,(j)} = \hat{\boldsymbol{\theta}}_h^{IV,(j)},$$

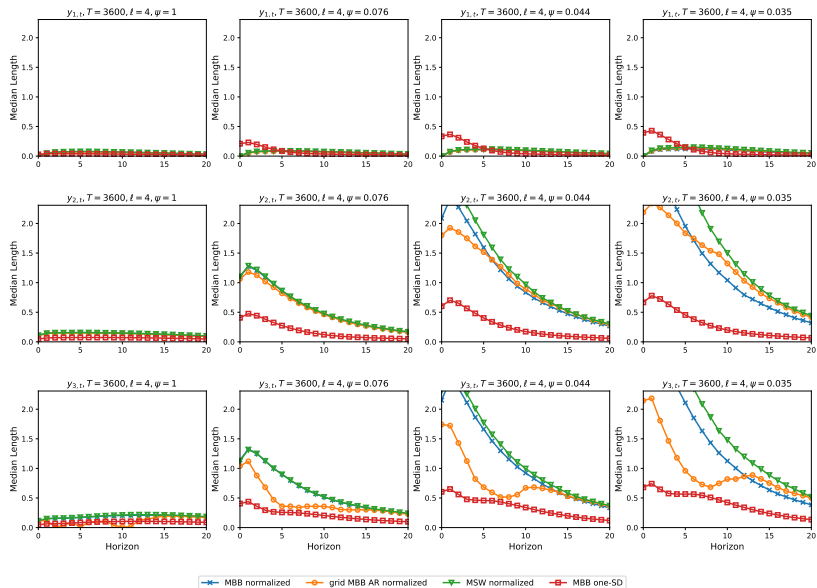
so that the same ratio estimator has two readings: an impact-normalized structural-shock IRF, or a causal effect of an observed exogenous movement in $y_{j,t}$.

This is not surprising: with the same lags, controls, and normalization, the LP-IV and SVAR-IV estimators coincide at impact, $h = 0$.

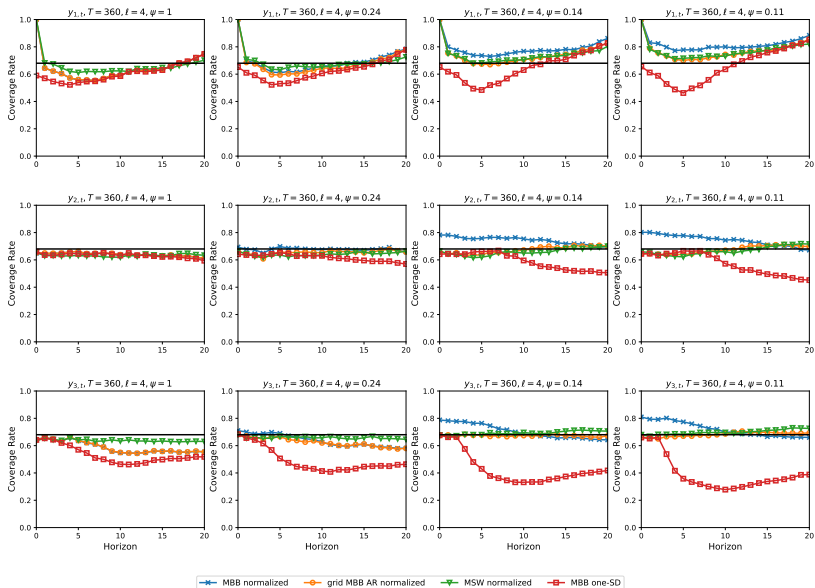
Appendix - full IRF coverage, $T = 3600$



Appendix - full IRF median length, $T = 3600$



Appendix - full IRF coverage, $T = 360$



Appendix - full IRF median length, $T = 360$

