

Lecture 4: Time Series Approaches to Identifying Macroeconomic Shocks

Christian Wolf

MIT

14.461, Fall 2022

Motivation

- We have seen: structural business-cycle models admit **SVMA representations**

$$y_t = \sum_{\ell=0}^{\infty} \Theta_{\ell} \varepsilon_{t-\ell}$$

From the Θ 's we get our objects of interest: **IRFs, FVDs, HDs**

- How can we use data to learn about the Θ 's? Three broad approaches:
 1. **Semi-structural, time series**: use time-series properties of y_t + identifying assumptions that hold for *families* of models **This will be our focus for the next few lectures.**
 2. **Semi-structural, cross sectional**: what can microeconomic causal designs teach us?
Will discuss this towards the end.
 3. **Structural**: specify full-blown model, then *estimate* this model using micro/macro data
Will discuss throughout to help when semi-structural approaches fall short.

Outline

1. The Time Series Identification Challenge

2. Invertibility + X

Exact Exclusion Restrictions

Sign Restrictions

Other Approaches

3. Identification without Invertibility

Instruments/Proxies

Recoverability

Outline

1. The Time Series Identification Challenge

2. Invertibility + X

Exact Exclusion Restrictions

Sign Restrictions

Other Approaches

3. Identification without Invertibility

Instruments/Proxies

Recoverability

Identifying SVMA models

- Motivated by our structural modeling results, we are willing to **assume** an **SVMA model** for our observed macro aggregates y_t :

$$y_t = \sum_{\ell=0}^{\infty} \Theta_{\ell} \varepsilon_{t-\ell}, \quad \varepsilon_t \sim WN(0, I)$$

Note that we may well have $n_{\varepsilon} > n_y$.

- **Identification challenge**

- Can estimate the **second-moment properties** of y_t (i.e., autocovariances/spectrum/Wold)
- The identification challenge—as we will see—is that infinitely **many VMA models** are consistent with those second-moment properties

[Of course at some level this is unsurprising. Why should simple covariances tell us something about the structural origins of business cycles or about the effects of policy interventions?]

The Wold decomposition

- Let's first of all show that we can always find **a VMA representation**
- We can do so easily using the **Wold decomposition** of y_t :

$$y_t = \tilde{\Psi}(L)\tilde{\varepsilon}_t$$

where $\tilde{\varepsilon}_t \equiv \Sigma^{-1/2}u_t \sim WN(0, I)$ and

$$u_t \equiv y_t - \mathbb{E}^*(y_t \mid \{y_\tau\}_{-\infty < \tau < t}), \quad u_t \sim WN(0, \Sigma)$$

- Clearly there is nothing “structural” about the $\tilde{\varepsilon}_t$'s—they are just **orthogonal innovations in a representation of y_t** , scaled to have unit variance
- Let's show that there exist many such **observationally equivalent** representations, none of which necessarily have any economic meaning ...

Indeterminacy I: rotations

- The first indeterminacy is what I'll refer to as **static indeterminacy**

[note that this is my language, not standard]

- Start with the true **SVMA model**, and assume that $n_\varepsilon = n_y = n$. Now let $Q \in O(n)$ denote an $n \times n$ **orthogonal matrix**. Then the process

$$y_t = \sum_{\ell=0}^{\infty} \tilde{\Theta}_\ell \tilde{\varepsilon}_{t-\ell}, \quad \tilde{\varepsilon}_t \sim WN(0, I)$$

with $\tilde{\Theta}_\ell = \Theta_\ell Q'$ and $\tilde{\varepsilon}_t = Q\varepsilon_t$ has the same **second-moment properties**:

$$\text{Cov}(y_t, y_{t-h}) = \sum_{\ell=0}^{\infty} \Theta_\ell \Theta'_{\ell+h} = \sum_{\ell=0}^{\infty} \tilde{\Theta}_\ell Q Q' \tilde{\Theta}'_{\ell+h}$$

- In words: **orthogonal rotations** of today's shocks leave covariances unchanged
 - This is a problem: we could just relabel the shocks or take linear combinations—the data wouldn't be able to reject it

Indeterminacy II: root flipping

- The second indeterminacy is what I'll refer to as **dynamic indeterminacy**
 - To illustrate this consider a simple univariate MA(2), i.e. $n_y = n_\varepsilon = 1$:

$$y_t = \theta_0 \varepsilon_t + \theta_1 \varepsilon_{t-1}$$

- Now compare this with the following alternative process:

$$y_t = \theta_1 \varepsilon_t + \theta_0 \varepsilon_{t-1}$$

Note how they share all second-moment properties! But of course the implied shock IRFs look very different as long as $\theta_0 \neq \theta_1$...[\[what are they?\]](#)

- What did we do? We took the polynomial $\Theta(L) = \theta_0 + \theta_1 L$ and “flipped the root”
- Lesson: **flipping roots** of MA polynomials leaves second moments unchanged
[For a general multivariate treatment of this see e.g. Lippi and Reichlin \(1994\)](#)

Indeterminacy III: system size

- The third determinacy is what I'll call **size indeterminacy**
 - So far we've assumed that $n_y = n_\varepsilon$. Our options for possible SVMA processes expand even more if we allow $n_\varepsilon > n_y$
 - To illustrate suppose again that the true model is the same one-shock MA(2) as before, but now we entertain the alternative *two*-shock process

$$y_t = \theta_{01}\varepsilon_{1t} + \theta_{02}\varepsilon_{2t} + \theta_{11}\varepsilon_{1,t-1} + \theta_{12}\varepsilon_{2,t-1}$$

- This process is just as consistent with the data as long as

$$\begin{aligned}\text{Var}(y_t) &= \theta_{01}^2 + \theta_{02}^2 + \theta_{11}^2 + \theta_{12}^2 = \theta_0^2 + \theta_1^2 \\ \text{Cov}(y_t, y_{t-1}) &= \theta_{01}\theta_{11} + \theta_{02}\theta_{12} = \theta_0\theta_1\end{aligned}$$

But of course it implies very different IRFs/FVDs/HDs ...

- Lesson: we can always split any given n_y -dimensional covariance structure into **many** ($n_\varepsilon > n_y$) **distinct shocks**

Summary: identification challenge

- From second moments alone, our SVMA model is severely **under-identified**
Aside: we will discuss the feasibility and desirability of identification from higher-order moments later
- At some level this is unsurprising—we so far haven't made any **economic identifying assumptions** that would allow claims about causality
- Rest of this lecture: progress with **as little structure as possible**
 1. **Invertibility + X**: zero, sign, statistical, ...
This is the classical “VAR” literature.
 2. **Identification without invertibility**: macro IVs, recoverability + X
The “macro IV” approach is very popular these days, and looks rather similar to standard microeconomic/“credibility revolution” practice.

Outline

1. The Time Series Identification Challenge

2. Invertibility + X

Exact Exclusion Restrictions

Sign Restrictions

Other Approaches

3. Identification without Invertibility

Instruments/Proxies

Recoverability

Invertibility

- Recall our definition of **invertibility**:

Definition

A VMA process $\{y_t\}$ is said to be **invertible** with respect to $\{\varepsilon_t\}$ if

$$\varepsilon_t \in \text{span}(y_\tau, -\infty < \tau \leq t)$$

- This is an economically **substantive assumption**:
 - It implies that an **econometrician observing $\{y_\tau\}_{-\infty < \tau < t}$** can predict just as well as an **agent living in the model and observing $\{\varepsilon_\tau\}_{-\infty < \tau < t}$**
 - Certainly need $n_\varepsilon \leq n_y$. The more (relevant) data we include, the more plausible this as'n
 - It's often not satisfied for standard macro models with standard observables, even if $n_\varepsilon = n_y$. Let's consider one very simple example. **Following Fernandez-Villaverde et al. (2007).**

Invertibility: illustration

- Consider a simple **permanent income consumption model**. Income is

$$e_t = \sigma_\varepsilon \varepsilon_t$$

and consumption is

$$c_t = c_{t-1} + \frac{r}{1+r} \sigma_\varepsilon \varepsilon_t$$

- Suppose we observe savings growth $y_t \equiv \Delta s_t$, where $s_t = e_t - c_t$:

$$y_t = \frac{1}{1+r} \sigma_\varepsilon \varepsilon_t - \sigma_\varepsilon \varepsilon_{t-1}$$

Is this model **invertible**?

- We have **one shock and one observable**, so things are looking good
- But unfortunately it doesn't work out, as there is no convergence in mean square:

$$\varepsilon_t = \frac{1}{\sigma_\varepsilon} \sum_{\ell=0}^T (1+r)^{\ell+1} y_{t-\ell} + (1+r)^{T+1} \varepsilon_{t-T}$$

Invertibility: discussion

- More general point: invertibility can fail **for many reasons**
 1. A necessary—but very strong—condition is that we have **as many observables as shocks**
 2. But even that is not sufficient. Key culprits for non-invertibility when $n_\varepsilon = n_y$ are:
 - a) There are **news shocks** (e.g., forward guidance, fiscal spending plans)
Like our MA from the previous slide—larger lagged than contemporaneous effect naturally generates non-invertibility. See Leeper et al. (2013) for a detailed discussion.
 - b) There are **noise shocks** (e.g., noise that is misperceived as news about future fundamentals)
See Chahrour-Jurado (2018) for a detailed discussion.

Intuition: with those shocks we would need to be able to look into the future. $n_y = n_\varepsilon$ in general only ensures that $\varepsilon_t \in \text{span}(y_\tau, -\infty < \tau < \infty)$, not that just looking to the *past* is enough. We will return to this when we discuss recoverability.

- So what do we gain from this strong assumption? As it turns out, a lot ...

Invertibility & Wold innovations

- **Invertibility** has important implications for **Wold innovations**. Recall that

$$u_t \equiv y_t - \mathbb{E}^*(y_t \mid \{y_\tau\}_{-\infty < \tau < t})$$

- Since y and u share the **same span**, it follows that, *under invertibility*,

$$\varepsilon_t \in \text{span}(u_\tau, -\infty < \tau \leq t)$$

But ε_t is uncorrelated with all u_τ for $\tau < t$, and u_t is white noise, so $\varepsilon_t \in \text{span}(u_t)$ or

$$H\varepsilon_t = u_t$$

for some matrix H , where $HH' = \Sigma = \text{Var}(u_t)$

- Also recall that the Wold innovations u_t and their “causal effects” $\Psi(L)$ are identifiable. We have thus made a lot of progress ...

Invertibility & Wold innovations

- By imposing invertibility, we have identified the structural shocks ε_t and their dynamic causal effects up to the **static rotation problem**:

1. The **structural shocks** are given as

$$\varepsilon_t = Q\Sigma^{-1/2}u_t$$

2. Their **dynamic causal effects** are given as

$$\Theta(L) = \Psi(L)\Sigma^{-1/2}Q'$$

- Invertibility has thus helped with the identification problem by eliminating the **dynamic & size indeterminacy**. Remains to find identifying assumptions to pin down Q ...
Intuitively: $n_y = n_\varepsilon$ (which is required for invertibility) rules out size indeterminacy, and then the sufficiency of *past* observables rules out the dynamic MA root flipping.

Connection to VARs

- Note that this approach to identification can be *operationalized* using **VARs**:
 - We have the SVMA model $y_t = \Theta(L)\varepsilon_t$. By invertibility, we know that the one-sided inverse $\tilde{A}(L) \equiv \Theta(L)^{-1}$ exists.
Strictly speaking, invertibility only rules out roots outside the unit circle, not on it. Will ignore.
 - Letting $A(L) = H^{-1}\tilde{A}(L)$, we thus have

$$A(L)y_t = u_t = H\varepsilon_t$$

That is, VAR residuals = Wold residuals, which we now just need to rotate.

- The derivations above are why the literature often speaks of “**VAR identification**”
 - I find this counterproductive. The identifying assumptions are **invertibility + something** to pick out the right notation. We will now discuss options for this “something else” ...

Outline

1. The Time Series Identification Challenge

2. Invertibility + X

Exact Exclusion Restrictions

Sign Restrictions

Other Approaches

3. Identification without Invertibility

Instruments/Proxies

Recoverability

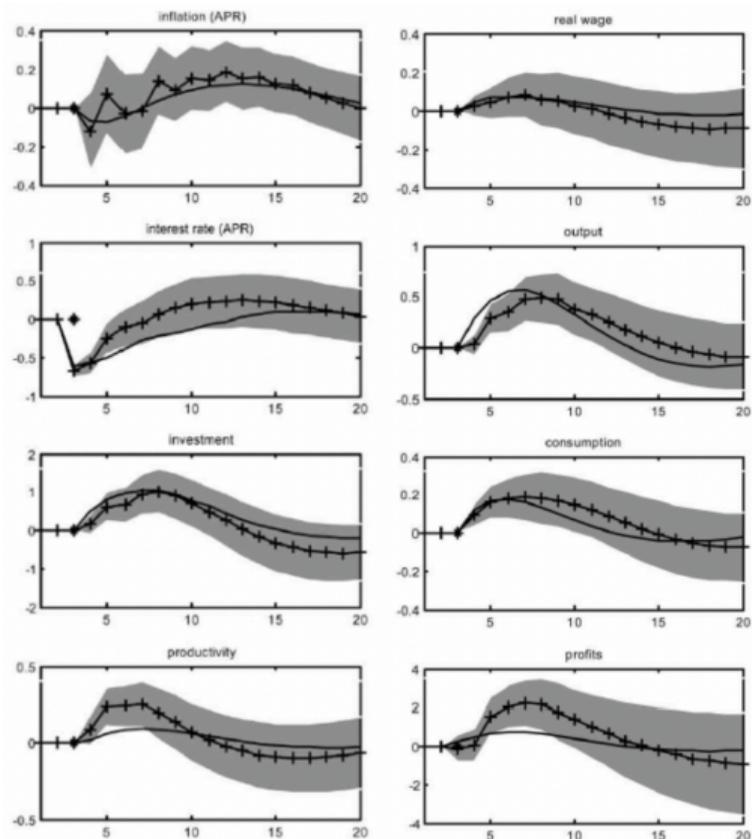
Exact exclusion restrictions

- Traditional approach: point-identify the system through *exact* zero restrictions on $\Theta(L)$
 - Need to impose enough restrictions such that only one $Q \in O(n)$ is consistent with them
- We will consider examples of the most well-known examples of such restrictions:
 1. Restrictions on **short-run impulse responses** Θ_0
We will review Christiano-Eichenbaum-Evans (1999, 2005).
 2. Restrictions on **long-run impulse responses** $\sum_{\ell=0}^{\infty} \Theta_{\ell}$
We will review Blanchard-Quah (1991).

Short-run identification

- Consider the assumption that Θ_0 is **lower-triangular**:
 - In economic terms: variable 1 only responds contemporaneously to shock 1, variable 2 responds contemporaneously to shocks 1 & 2, and so on
 - Mathematically: $H = \text{chol}(\Sigma)$
- **Christiano et al. (1999, 2005)** operationalize this as 'n to identify **monetary policy shocks**
 - Observables y_t : real GDP, real consumption, GDP deflator, real investment, real wage, labor productivity, federal funds rate, real profits, growth rate of M2
 - Assumptions: invertibility + recursive shock ordering
Here this means that all variables listed before the fed funds rate don't respond within the period to monetary shocks, while monetary policy doesn't respond to changes in the variables listed after.
- This identification approach gave the “canonical” monetary policy IRFs [*next slide*]
Note that these IRFs look materially different on recent samples, e.g. see Barakchian-Crowe (2013)

Short-run identification



Can you think of any problems with this?

Long-run identification

- Suppose we are willing to impose that the **long-run response** of variable i to shock j is zero. That implies:

$$0 = \sum_{\ell=0}^{\infty} \Theta_{i,j,\ell} = \sum_{\ell=0}^{\infty} \Psi_{i,\bullet,\ell} H_{\bullet,j} = \sum_{\ell=0}^{\infty} \Psi_{i,\bullet,\ell} \Sigma^{1/2} Q'_{\bullet,j} \quad (1)$$

- (1) thus imposes an additional restriction on the orthogonal matrix Q , thus aiding with identification of the system
- **Blanchard-Quah (1991)** operationalize this as'n to disentangle **demand & supply shocks**
 - Observables y_t : real output, unemployment
 - Assumptions: invertibility + demand shocks do not affect real output

Long-run identification

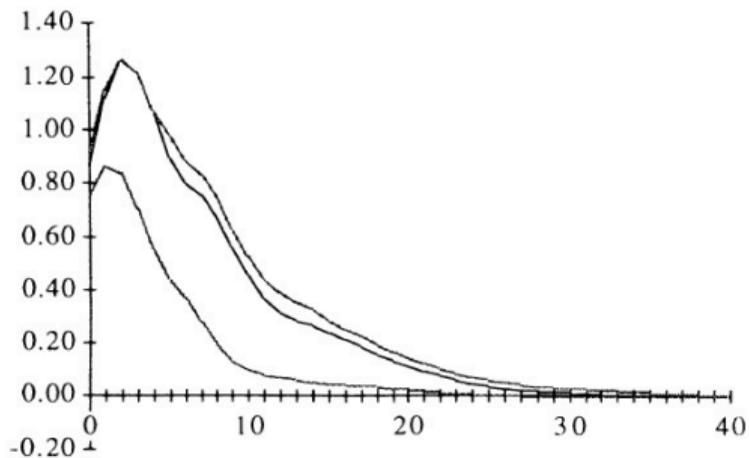


FIGURE 3. OUTPUT RESPONSE TO DEMAND

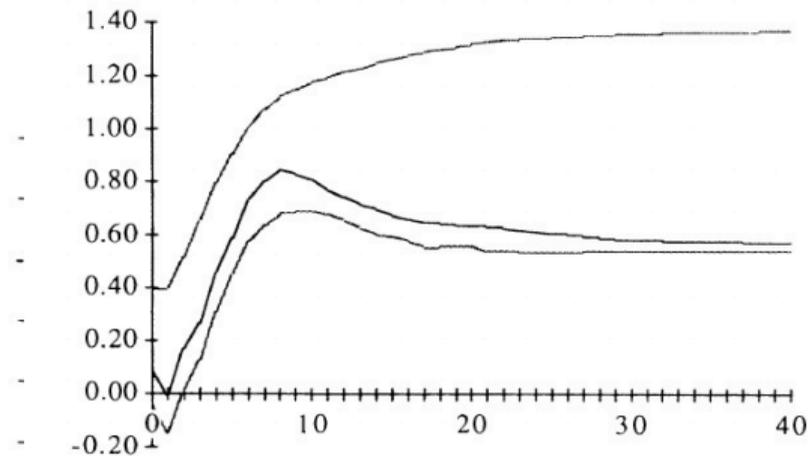


FIGURE 4. OUTPUT RESPONSE TO SUPPLY

Can you think of any problems with this?

Outline

1. The Time Series Identification Challenge

2. Invertibility + X

Exact Exclusion Restrictions

Sign Restrictions

Other Approaches

3. Identification without Invertibility

Instruments/Proxies

Recoverability

Motivation

- Exact zero restrictions are widely regarded as implausible
E.g.: effects of monetary policy on output may be delayed, not exactly 0 on impact.
- One popular alternative: impose restrictions on **signs** of impulse responses
 - Example: a contractionary monetary shock should probably lead to an **increase in interest rates** and a **decrease in prices** (eventually)
 - This is promising: if central banks lean against inflation, then all non-policy shocks lead to a *positive co-movement of output and inflation*
- **Q**: How can we operationalize this intuition? What can invertibility + sign restrictions tell us about the **SVMA model**?

Invertibility + sign restrictions

- Recall that IRFs for a given rotation Q are equal to

$$\Theta(L) = \Psi(L)H = \Psi(L)\Sigma^{1/2}Q'$$

- Suppose we want to impose that $\Theta_{i,j,\ell} \geq 0$. This rules out some Q 's! We need

$$\Psi_{i,\bullet,\ell}\Sigma^{1/2}Q'_{\bullet,j} \geq 0 \quad (2)$$

\Rightarrow Sign restrictions give **identified sets**: keep all Q 's such that (i) $QQ' = I$ and (ii) all imposed sign restrictions of the form (2) hold

Example: sign restrictions for monetary policy

- Let's illustrate the workings of sign restrictions through a **toy model example**:

$$y_t = \mathbb{E}_t(y_{t+1}) - (i_t - \mathbb{E}_t(\pi_{t+1})) + \sigma^d \varepsilon_t^d \quad (\text{IS})$$

$$\pi_t = \kappa y_t + \beta \mathbb{E}_t(\pi_{t+1}) - \sigma^s \varepsilon_t^s \quad (\text{NKPC})$$

$$i_t = \phi_\pi \pi_t + \sigma^m \varepsilon_t^m \quad (\text{TR})$$

$$\varepsilon_t \equiv (\varepsilon_t^d, \varepsilon_t^s, \varepsilon_t^m) \sim WN(0, I)$$

- Solving the model gives a **static mapping** from shocks to observables: [i.e., $\text{VMA}(0)$]

$$\begin{pmatrix} y_t \\ \pi_t \\ i_t \end{pmatrix} = \underbrace{\begin{pmatrix} + & + & - \\ + & - & - \\ + & - & + \end{pmatrix}}_{\Theta} \times \begin{pmatrix} \varepsilon_t^d \\ \varepsilon_t^s \\ \varepsilon_t^m \end{pmatrix}$$

Denote the response of variable j to shock k by Θ_{jk} .

Example: sign restrictions for monetary policy

- Interpreting the Q 's

- Let $x_t = (y_t, \pi_t, i_t)'$ and $\Sigma_x = \text{Var}(x_t)$. Note that

$$\Sigma_x \equiv \text{Var}(x_t) = \Theta\Theta'$$

- Before we defined the Q 's relative $\Sigma^{1/2}$. More generally we can define them w.r.t. any matrix A such that $AA' = \Sigma_x$, including Θ . The Q 's are just **rotations w.r.t. a basis**.
- Choosing Θ as the basis has an advantage: if $\tilde{\Theta} = \Theta Q'$, then the matrix Q corresponds to **mis-identified structural shocks $\tilde{\varepsilon}_t$** given as

$$\tilde{\varepsilon}_t \equiv Q\varepsilon_t$$

Of course you can only do this w.r.t. an underlying model, not in any actual empirical application. Our point here is illustration of the workings of sign restrictions.

- We'd like to find identifying assumptions that are consistent with $Q = I$ (so $\tilde{\varepsilon}_t = \varepsilon_t$) but rule out (almost) everything else. **Q**: will **sign restrictions** on Θ do the trick?

Example: sign restrictions for monetary policy

- Let's impose the following **minimalist sign restrictions**:

$$\Theta = \begin{pmatrix} ? & ? & ? \\ ? & ? & - \\ ? & ? & + \end{pmatrix} \quad (3)$$

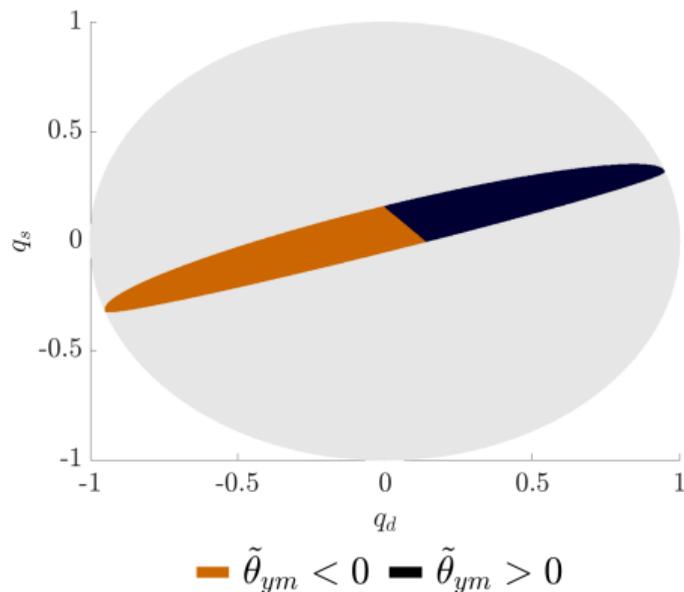
- Is this enough to learn anything about the **real effects of monetary policy**, Θ_{ym} ?
 - Promising: monetary shocks are the *only* shocks to move π_t and i_t in **opposite directions**
 - This implies that $\tilde{\varepsilon}_t^m = \varepsilon_t^m$ is consistent with (3), while $\tilde{\varepsilon}_t^m = \varepsilon_t^d$ or $\tilde{\varepsilon}_t^m = \varepsilon_t^s$ are not
- Formally: (3) gives us an *identified set* $[\underline{\Theta}_{ym}, \overline{\Theta}_{ym}]$ for Θ_{ym} , defined via the programs

$$\inf_q / \sup_q \quad \Theta_{1,\bullet} \cdot q$$

$$\text{s.t. } \|q\| = 1 \text{ and } \Theta_{2,\bullet} \cdot q < 0, \Theta_{3,\bullet} \cdot q > 0$$

The identified set & masquerading shocks

What does the identified set $[\underline{\Theta}_{ym}, \bar{\Theta}_{ym}]$ look like?



- Since Θ is invertible, it's easy to see that we can't sign Θ_{ym} : [proof on board]

$$\underline{\Theta}_{ym} < 0 < \bar{\Theta}_{ym}$$

- Interpretation: “masquerading” shocks
 - The truth ($q = (0, 0, 1)'$) is in the identified set, while $q = (1, 0, 0)'$ and $q = (0, 1, 0)'$ are not
 - But: linear combos of expansionary supply and demand shocks can also move $i \uparrow$ & $\pi \downarrow$, but $y \uparrow$
- In general, the identified set for Q may be empty, contain a singleton, or contain multiple elements. The last case (as seen here) is typical.

Econometric aside: the Haar prior

- Much of the (early) sign restrictions literature **doesn't report identified sets**
- Instead: additionally impose a **prior on orthogonal rotation matrices**, and then characterize the implied posterior
 - Common choice (for computational reasons): **"uniform" Haar prior**
 - The effect of this prior is to **re-weight the parameters in the identified set**. Desirable?
- Briefly: no. The prior is **funky in economic terms** and can be central to posterior tightness in actual applications. **See Baumeister-Hamilton (2015) and Wolf (2020)**.
 - Best practice: frequentist inference on entire identified set or Bayesian prior-robust inference **Moon-Schorfheide (2012), Giacomini-Kitagawa (2021)**

Outline

1. The Time Series Identification Challenge

2. Invertibility + X

Exact Exclusion Restrictions

Sign Restrictions

Other Approaches

3. Identification without Invertibility

Instruments/Proxies

Recoverability

What other approaches for **identification under invertibility** are there?

1. Combining **zero and sign restrictions** [Arias, Rubio-Ramírez and Waggoner \(2018\)](#)
 - Can yield much tighter identified sets than sign restrictions alone, even without additional imposition of the Haar prior
2. Sign/magnitude restrictions on **other objects**, not IRFs. Examples:
 - Shock that explains a large share of fluctuations in variable X [Angeletos et al. \(2020\)](#)
 - Shock that explains most volatility at a particular date [Antolín-Díaz & Rubio-Ramírez \(2018\)](#)
3. Explicitly defended **probabilistic priors on Q** (rather than the Haar prior)
[Baumeister and Hamilton \(2015\)](#), [Plagborg-Møller \(2019\)](#)
4. Use more information than just second moments (= **“statistical” identification**)

Statistical identification

- Identification through **higher-order moments** has recently gained in popularity
- Basic idea:
 - Lack of identification so far because we only looked at second moments. If shocks are **not normal**, higher-order moments can help
 1. Next slide: identification through heteroskedasticity point-identifies the model
To me the most promising of these “statistical” approaches, since there’s still a lot of economics in trying to come up with reasons for why relative shock volatilities change.
 2. Alternative: if true shocks are **independent and non-normal**, then also point-identified
Intuition: only the true rotation of the u_t ’s is independent. See Gouriéroux et al. (2017).
 - But note: even with statistical id, you still need economic reasoning to *label* the shocks
E.g.: which of the “statistically” identified shocks is the monetary policy shock?

Identification through heteroskedasticity in one slide

- Suppose now that the struct. shocks ε_t have different volatilities on two subsamples
 - Write the variances as $\{(\sigma_j^a)^2, (\sigma_j^b)^2\}_{j=1}^n$ for our two subsamples a and b **Note: now we are not normalizing the shock variances, so we instead normalize the diagonal entries of H to 1.**
 - For each sub-sample $k = a, b$, the Wold residuals u_t^k satisfy

$$\text{Var}(u_t^k) = H \begin{pmatrix} (\sigma_1^k)^2 & \dots & 0 \\ & \ddots & \\ 0 & \dots & (\sigma_n^k)^2 \end{pmatrix} H' = \Sigma_k$$

- We thus now have more information that we can exploit:

$$\Sigma_b \Sigma_a^{-1} = H \begin{pmatrix} (\frac{\sigma_1^b}{\sigma_1^a})^2 & \dots & 0 \\ & \ddots & \\ 0 & \dots & (\frac{\sigma_n^b}{\sigma_n^a})^2 \end{pmatrix} H^{-1}$$

- Columns of H are eigenvectors of $\Sigma_b \Sigma_a^{-1}$, which is identified. Given our normalization of the diagonal of H to 1, we see that H is identified if the volatility ratios differ across shocks

Graphical intuition

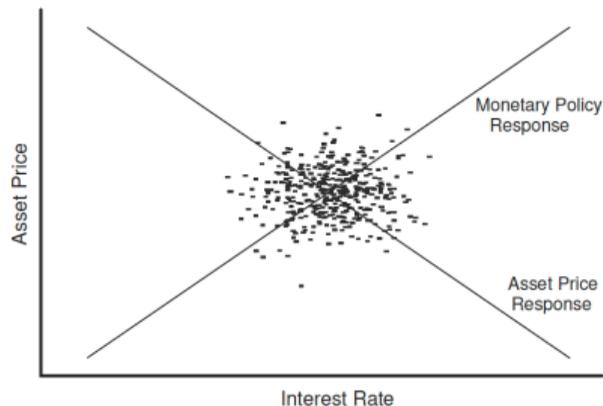


Fig. 1. Joint determination of interest rates and asset prices.

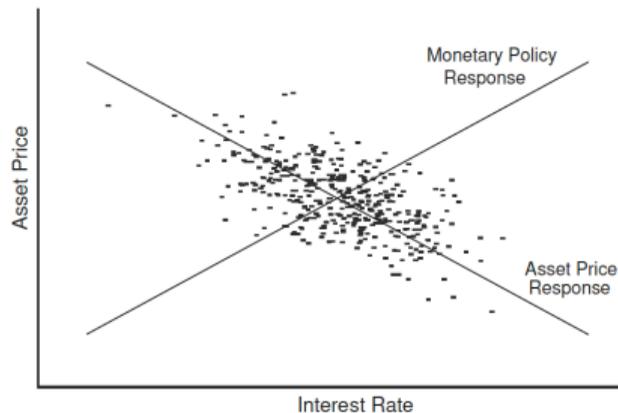


Fig. 2. Policy dates.

- Key idea: only two slopes can jointly rationalize these two clouds
 - Easiest to see if one shock is *dominant* in one regime = event study. But not necessary.
 - Of course then still need to label the two lines = two shocks
Here: what is causal effect of MP and what is rule response? Other natural application: what's demand and what's supply?

Some thoughts on statistical identification

- Comparison with our **standard identification arguments**
 - Classical invertibility-based analysis assumes **white noise shocks**, e.g. allowing for stochastic volatility. Just doesn't **exploit** higher-order moments
 - Statistical approaches achieve identification by strengthening assumptions on the shock process: independence/conditional orthogonality rather than just white noise
 - Some **words of caution**: See Montiel-Olea et al. (2022) for details.
 - For heteroskedasticity: must assume that only **volatility ratios** changed, but not **shock propagation**. Also can't handle **common volatility changes** (Great Moderation?).
 - Higher-order moments are **hard to estimate**, in particular in time series. Analysis is necessarily fragile w.r.t. statistical properties of shocks, e.g. **weak identification** issues arise if the shocks are nearly i.i.d. Gaussian
- ⇒ To me reliance on **economic identifying asns** is a virtue, not a bug.

Outline

1. The Time Series Identification Challenge

2. Invertibility + X

Exact Exclusion Restrictions

Sign Restrictions

Other Approaches

3. Identification without Invertibility

Instruments/Proxies

Recoverability

Identification without invertibility

- Recently there's been a push towards **invertibility-robust** identification approaches
- To me the most promising of these is the **narrative/IV** approach

- Assumes that, through institutional knowledge, the researcher can find a valid IV for a macro shock:

$$z_t = \alpha \varepsilon_{1,t} + \sigma_\nu \nu_t, \quad \nu_t \sim WN(0, 1) \quad (4)$$

- This brings **macroeconometrics closer to microeconomic practice**: try to find credible natural experiments that satisfy (4). We will thus focus on such IVs.
 - Main appeal: requires **no** further assumptions on the **SVMA model**, at least for IRFs
- Content of the next couple of slides:
 0. Clarify what goes wrong in “VAR” analysis **without invertibility**
 1. Discuss identification based on **IVs** (4) (including some popular examples)
 2. Briefly: sketch **“recoverability”-based identification**

Degree of invertibility

- Suppose we wish to use invertibility-based identification to study a **single shock** $\varepsilon_{j,t}$. This can work if and only if that shock is invertible:

$$\varepsilon_{j,t} \in \text{span}(y_\tau, -\infty < \tau \leq t) \Leftrightarrow \varepsilon_{j,t} = h' u_t \text{ for some } h$$

- This condition is weaker than invertibility of all shocks. It's commonly referred to as partial invertibility. [Forni-Gambetti-Sala \(2019\)](#)
- Invertibility shouldn't be an either-or proposition. Intuitively, we should be able to do well if $\varepsilon_{1,t}$ is "**close to invertible**". How to formalize that?

$$\mathcal{R}_{j,0}^2 \equiv 1 - \frac{\text{Var}^*(\varepsilon_{j,t} \mid \{y_\tau\}_{-\infty < \tau \leq t})}{\text{Var}(\varepsilon_{j,t})} = 1 - \text{Var}^*(\varepsilon_{j,t} \mid \{y_\tau\}_{-\infty < \tau \leq t})$$

- Can show: the asymptotic bias of many VAR-type procedures is a function of $\mathcal{R}_{j,0}^2$, vanishing as $\mathcal{R}_{j,0}^2 \rightarrow 1$ [Forni-Gambetti-Sala \(2019\)](#), [Plagborg-Møller & Wolf \(2021\)](#)
- Appendix: what do we in general get in the non-invertible case? [Details](#)

Outline

1. The Time Series Identification Challenge

2. Invertibility + X

Exact Exclusion Restrictions

Sign Restrictions

Other Approaches

3. Identification without Invertibility

Instruments/Proxies

Recoverability

Setting & identification problem

- Let's suppose we have access to a **macro IV/proxy**:

$$z_t = \alpha \varepsilon_{1,t} + \sigma_\nu \nu_t, \quad \nu_t \sim WN(0, 1)$$

- How may we be able to find such z 's?
 - **Institutional knowledge** of policy decisions (e.g. fiscal spending \perp to macro conditions)
 - **High-frequency** financial market responses (e.g. idiosyncratic global oil supply changes)
- Even if found, such z 's are unlikely to capture *all* of the hidden shock $\varepsilon_{1,t}$, so we'll invariably have measurement error ν_t
- **Q**: how can we use such IVs to identify our objects of interest?
 - Note that the econometric model is now the **general SVMA** + our **IV equation**:

$$\begin{aligned} y_t &= \Theta(L)\varepsilon_t, \\ z_t &= \alpha \varepsilon_{1,t} + \sigma_\nu \nu_t, \quad (\varepsilon_t', \nu_t)' \sim WN(0, I) \end{aligned}$$

Dynamic causal effects

- It's trivial to identify **relative dynamic causal effects**

- Note that we have

$$\text{Cov}(y_{i,t+\ell}, z_t) = \Theta_{i,1,\ell} \text{Cov}(\varepsilon_{1,t}, \alpha\varepsilon_{1,t} + \sigma_\nu \nu_t) = \alpha\Theta_{i,1,\ell}$$

- Thus we can point-identify the IRF *ratio*

$$\frac{\Theta_{i,1,\ell}}{\Theta_{1,1,0}} = \frac{\text{Cov}(y_{i,t+\ell}, z_t)}{\text{Cov}(y_{1,t}, z_t)}$$

- This is often all we want: how does output (y_i) respond to a monetary shock that increases nominal rates by 100bp on impact (y_1)?

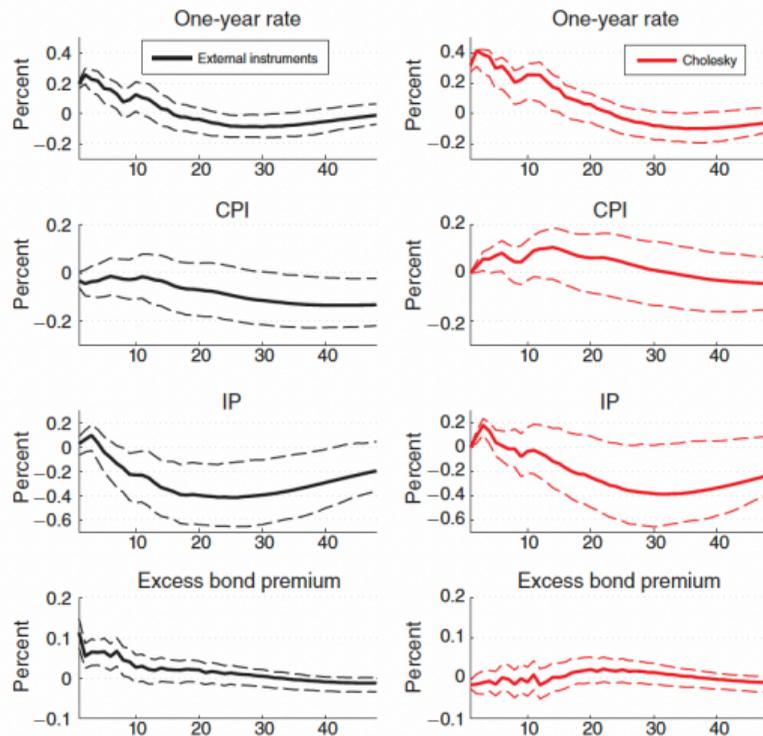
Absolute IRFs are only identified up to α , but we do not care.

- Note that none of these derivations assumed anything about **invertibility** (or more generally, our three sources of indeterminacy). Simply not needed for IV.

Examples of macro IVs

- Can we in practice hope to find such **macro IVs**?
- I will review two recent **canonical examples**
 1. High-frequency monetary policy surprises [Gertler-Karadi (2015)]
 - Construction: measure unexpected movements in interest rate futures around FOMC meeting/announcement dates
 - Identifying assumption: everyone knows rule $i_t = f(\Omega_t) + \varepsilon_t^m$ and information Ω_t , so surprise change must reflect ε_t^m [Can you think of problems with that?]
 2. High-frequency oil supply news [Känzig (2021)]
 - Construction: measure unexpected movements in oil futures prices around OPEC meeting/announcement dates
 - Identifying assumption: surprise responses reflect *only* news about future oil supply

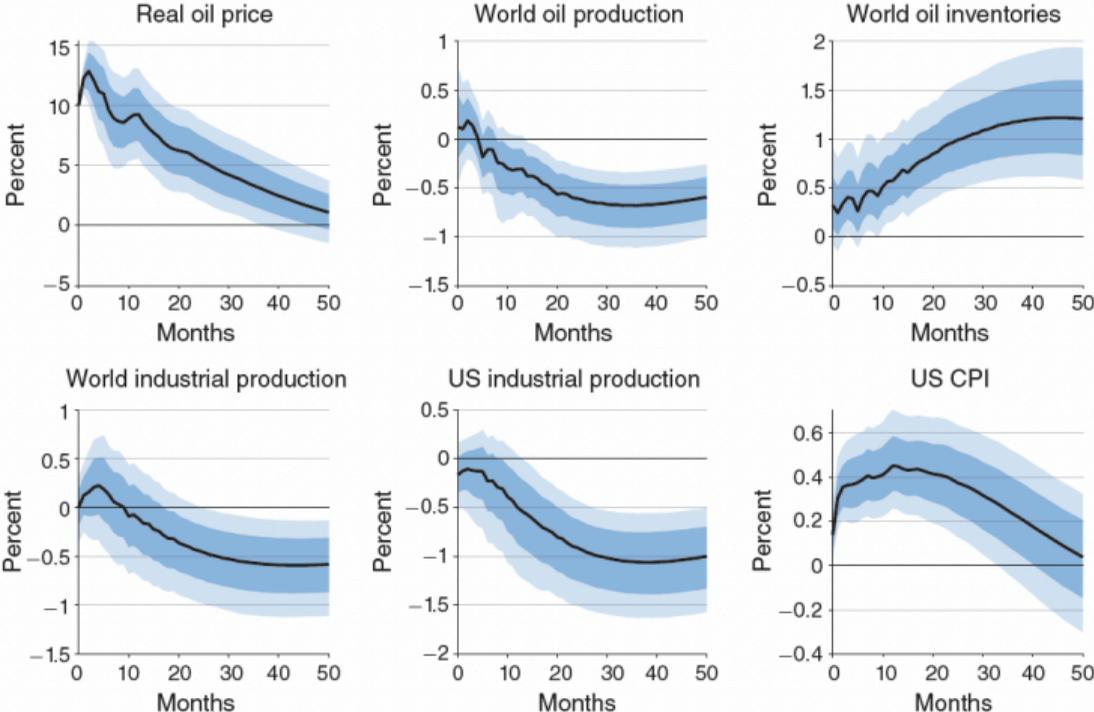
High-frequency monetary shocks



First-stage regression:
F: 21.55; Robust F: 17.64; R^2 : 7.76 percent; Adjusted R^2 : 7.40 percent

FIGURE 1. ONE-YEAR RATE SHOCK WITH EXCESS BOND PREMIUM

High-frequency oil shocks



First-stage regression: $F: 22.67$, robust $F: 10.55$, $R^2: 4.22\%$, Adjusted $R^2: 4.04\%$

FIGURE 3. IMPULSE RESPONSES TO AN OIL SUPPLY NEWS SHOCK

Aside: shock importance

- What can we say about **shock importance**, e.g. variance decompositions?

Issue: for this clearly need absolute impulse responses, not relative ones.

- Intuition on the **identification challenge**

- IV consists of signal ($\alpha\varepsilon_{1,t}$) and noise ($\sigma_\nu\nu_t$). We don't know the **signal-to-noise ratio**.
- This matters: e.g. can't know whether $\text{Corr}(y_{it}, z_t)$ is low because $\varepsilon_{1,t}$ is unimportant or because there's just a lot of noise
- Formally, our challenge is to say something about the **forecast variance ratio**:

$$\begin{aligned}\text{FVR}_{i,1,h} &\equiv 1 - \frac{\text{Var}(y_{i,t+h} \mid \{y_{t-\ell}\}_{\ell=0}^{\infty}, \{\varepsilon_{1,t+\ell}\}_{\ell=1}^{\infty})}{\text{Var}(y_{i,t+h} \mid \{y_{t-\ell}\}_{\ell=0}^{\infty})} \\ &= \frac{\sum_{m=0}^{h-1} \frac{1}{\alpha^2} \text{Cov}(y_{i,t+h}, z_t)^2}{\text{Var}(y_{i,t+h} \mid \{y_{t-\ell}\}_{\ell=0}^{\infty})}\end{aligned}$$

- Plagborg-Møller and Wolf (2021) prove that α and so FVR is **interval-identified**

[▶ Details](#)

IVs & invertibility

- Our IV analysis so far has proceeded without assuming **invertibility**—causal effects and shock importance are still (set-)identified without
- This is quite different from **early time series work** using macro IVs
 - The popular “**SVAR-IV**” approach uses IVs as a way to solve the rotation problem in invertibility-based identification [This is actually what's done in Gertler-Karadi.]
 - Suppose invertibility holds, so $u_t = H\varepsilon_t$. We can then recover $\varepsilon_{1,t}$ as the **projection of z_t on u_t** (which is $\propto \varepsilon_{1,t}$), rescaled to have unit variance:

$$\varepsilon_{1,t} = [\text{Cov}(z_t, u_t) \text{Var}(u_t)^{-1} \text{Cov}(z_t, u_t)']^{-1/2} \underbrace{\text{Cov}(z_t, u_t) \text{Var}(u_t)^{-1} u_t}_{\propto \varepsilon_{1,t}}$$

- Can then recover IRFs, variance decompositions, ... in the usual invertibility-based way
- But this all proceeded assuming invertibility, which we've seen is not needed. In fact IVs allow us to **test invertibility** ...

Testing invertibility

- The key insight is that, under invertibility, we must have

$$\mathbb{E}^*(y_{i,t} \mid \{z_\tau, y_\tau\}_{-\infty < \tau < t}) = \mathbb{E}^*(y_{i,t} \mid \{y_\tau\}_{-\infty < \tau < t})$$

- In words: z does not help to predict future y 's above and beyond the info contained in y
 - Why? $z_t = \alpha \varepsilon_{1,t} + \sigma_\nu \nu_t$, where the second term is useless for forecasting future y , and the first term is captured by past y 's
 - Suggests a simple test: does z Granger-cause y ?
As it turns out this is in fact the only testable implication of invertibility.
- If SVAR-IV is used even though **invertibility fails**, then we are subject to the issues reviewed in the slide appendix—bias that is a function of $\mathcal{R}_{j,0}^2$
 - Big picture: SVAR-IV is a bit of a “historical accident”, reflecting path dependence
 - Literature thought in terms of identification through invertibility (“SVAR identification”), so just viewing instruments as a way of finding the right Q seemed natural

Outline

1. The Time Series Identification Challenge

2. Invertibility + X

Exact Exclusion Restrictions

Sign Restrictions

Other Approaches

3. Identification without Invertibility

Instruments/Proxies

Recoverability

- Alternative approach: **recoverability-based identification**

Definition

A VARMA process $\{y_t\}$ is said to be **recoverable** with respect to $\{\varepsilon_t\}$ if

$$\varepsilon_t \in \text{span}(y_\tau, -\infty < \tau < \infty)$$

- Thus the identification problem can be solved by looking into the **past & future**:

$$\varepsilon_t = \sum_{\ell=-\infty}^{\infty} \Psi_\ell y_{t-\ell} = \sum_{\ell=0}^{\infty} Q_\ell u_{t+\ell} = Q_0 u_t + Q_1 u_{t+1} + \dots$$

- $\Psi(L)$ is a *two-sided* lag polynomial and $I = \sum_{\ell=0}^{\infty} Q_\ell \Sigma Q'_\ell$, $0 = \sum_{\ell=0}^{\infty} Q_\ell \Sigma Q'_{h+\ell} \quad \forall h > 0$
- The ID problem is now even harder: we are willing to call some n_y -dim. VMA “structural” (= assume away size indeterminacy), but still need to contend with static & *now dynamic* indeterminacy E.g.: Lippi-Reichlin (1994), Mertens-Ravn (2010), Chahrour-Jurado (2021)

- **This lecture note:** time series methods for identifying the SVMA Θ 's
 - We began by characterizing the **identification problem**: indeterminacy from three sources
 - We then reviewed the **solutions**: invertibility + X, macro IVs, recoverability + X
- **Next:** discuss various popular econometric strategies to in practice **implement these identification approaches** & so **estimate** the Θ 's

Appendix

Mis-identification without invertibility

- **Q:** what does invertibility-based identification recover without invertibility?
- We will derive everything from a **state-space representation**:

$$s_t = As_{t-1} + B\varepsilon_t \quad (5)$$

$$y_t = Cs_{t-1} + D\varepsilon_t \quad (6)$$

- Invertibility is about the informativeness of y_t about s_t (and so ε_t)
- We will compute this using the **Kalman filter**. Will rely on **linear projections**. Notation:

$$\hat{s}_{t|t} \equiv \mathbb{E}^* [s_t | \{y_\tau\}_{-\infty < \tau \leq t}]$$

$$\Sigma_{t|t}^s \equiv \text{var}^* [s_t | \{y_\tau\}_{-\infty < \tau \leq t}]$$

and similarly for y_t

- Algorithm: use standard linear projection formulas to update $\hat{s}_{t-1|t-1}$ and $\Sigma_{t-1|t-1}^s$ to $\hat{s}_{t|t}$ and $\Sigma_{t|t}^s$. Fixed point will give population limits (i.e., $t \rightarrow \infty$)

Kalman filtering algorithm

- Start with $\hat{s}_{t-1|t-1}$ and $\Sigma_{t-1|t-1}^s$. Predicting one period ahead:

$$\hat{s}_{t|t-1} = A\hat{s}_{t-1|t-1}$$

$$\Sigma_{t|t-1}^s = A\Sigma_{t-1|t-1}^s A' + BB'$$

$$\hat{y}_{t|t-1} = C\hat{s}_{t-1|t-1}$$

$$\Sigma_{t|t-1}^y = C\Sigma_{t-1|t-1}^s C' + DD'$$

- Next we use y_t :

$$\begin{aligned}\hat{s}_{t|t} &= \hat{s}_{t|t-1} + \mathbb{E}[(s_t - \hat{s}_{t|t-1})(y_t - \hat{y}_{t|t-1})'] \times \\ &\quad \mathbb{E}[(y_t - \hat{y}_{t|t-1})(y_t - \hat{y}_{t|t-1})']^{-1} (y_t - \hat{y}_{t|t-1})\end{aligned}$$

Kalman filtering algorithm

- Plugging in we get

$$\hat{s}_{t|t} = \hat{s}_{t|t-1} + \underbrace{\left[A\Sigma_{t-1|t-1}^s C' + BD' \right] \left[C\Sigma_{t-1|t-1}^s C' + DD' \right]^{-1}}_{\text{Kalman gain } K_t} (y_t - \hat{y}_{t|t-1})$$

Similarly

$$\Sigma_{t|t}^s = \Sigma_{t|t-1}^s - \left[A\Sigma_{t-1|t-1}^s C' + BD' \right] \left[C\Sigma_{t-1|t-1}^s C' + DD' \right]^{-1} \left[A\Sigma_{t-1|t-1}^s C' + BD' \right]'$$

- Thus, given any history of the observables $\{y_t\}_{t=0}^T$, we can construct a sequence of estimates of the hidden states s_t

Innovations representation

- Letting $t \rightarrow \infty$, we obtain the steady-state Kalman gain and state uncertainty:

$$\begin{aligned}\Sigma^s &= (A - KC)\Sigma^s(A - KC)' + BB' + KDD'K' - BD'K' - KDB' \\ K &= (A\Sigma^sC' + BD')(C\Sigma^sC' + DD')^{-1}\end{aligned}$$

- This allows us to re-write the state-space system in **innovations form**:

$$\hat{s}_t = A\hat{s}_{t-1} + K \underbrace{(y_t - \hat{y}_{t|t-1})}_{u_t}$$

$$y_t = C\hat{s}_{t-1} + u_t$$

with $\Sigma^u = C\Sigma^sC' + DD'$

- Now we finally get to the payoff: write **Wold innovations** in terms of the ε_t 's ...

Shocks and innovations

- The innovations representation gives

$$\begin{pmatrix} s_t \\ \hat{s}_t \end{pmatrix} = \begin{pmatrix} A & 0 \\ KC & A - KC \end{pmatrix} \begin{pmatrix} s_{t-1} \\ \hat{s}_{t-1} \end{pmatrix} + \begin{pmatrix} B \\ KD \end{pmatrix} \varepsilon_t$$
$$u_t = (C \quad -C) \begin{pmatrix} s_{t-1} \\ \hat{s}_{t-1} \end{pmatrix} + D\varepsilon_t$$

to arrive at

$$u_t = \left\{ D + (C \quad -C) \left(I - \begin{pmatrix} A & 0 \\ KC & A - KC \end{pmatrix} L \right)^{-1} \begin{pmatrix} B \\ KD \end{pmatrix} L \right\} \varepsilon_t = \sum_{\ell=0}^{\infty} M_{\ell} \varepsilon_t$$

- Under invertibility, $s_t = \hat{s}_t$, so $u_t = D\varepsilon_t$, exactly as we have seen
- Without invertibility, the identified shocks $\tilde{\varepsilon}_t = H^{-1}u_t$ are a **linear combination of current and past true shocks** ε_t
- Can show that the weight on $\varepsilon_{j,t}$ is bounded above by $\sqrt{\mathcal{R}_{j,0}^2}$, and that SVAR-IV attains this bound **Wolf (2020)**

Lower bound

- Finding a **lower bound** on shock importance is trivial
 - We must clearly have $\alpha^2 \leq \sigma_z^2$ – the IV can't be better than perfect
 - This upper bound on α gives a lower bound on the FVR, corresponding to naive regression of $y_{i,t+h}$ on z_t , ignoring measurement error
- In practice, IVs are noisy, so this lower bound is often **close to 0**

▶ back

Upper bound

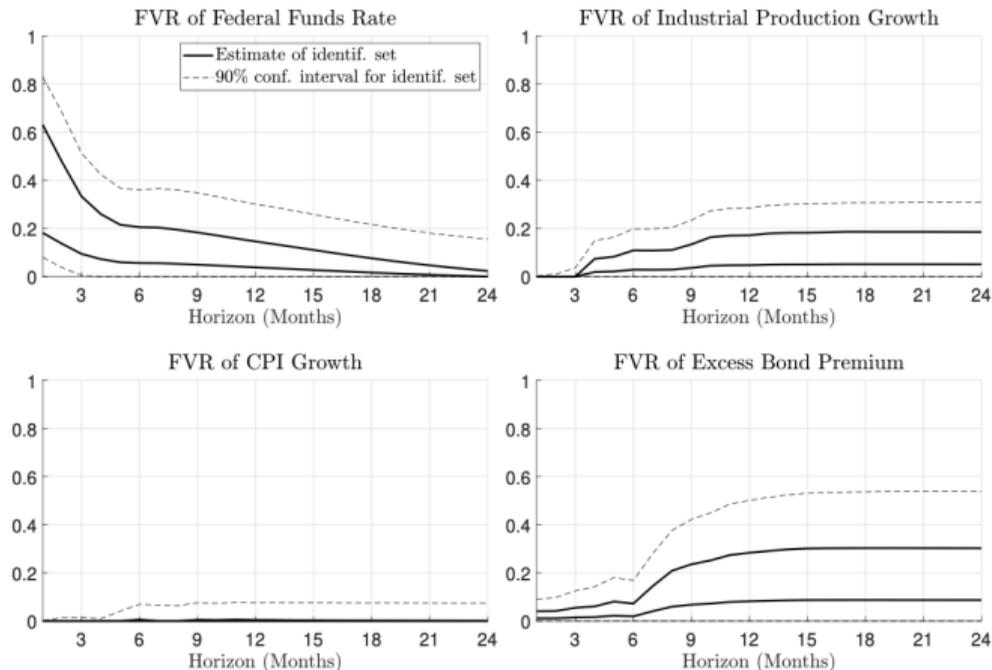
- Why should we expect there to be an **upper bound**? Example:
 - Suppose $y_{1,t}$ and $y_{2,t}$ have the same variance, but $\text{Cov}(y_{1,t}, z_t) \gg \text{Cov}(y_{2,t}, z_t)$. Then $\varepsilon_{1,t}$ can't explain much of $y_{2,t}$.
 - Why? If it explained much of $y_{2,t}$, it would *more than all of* $y_{1,t}$, which is impossible
- Let's formalize this intuition in a **static SVMA model**
 - The model now is $y_t = \Theta_0 \varepsilon_t$, and we let $\Sigma_y = \text{Var}(y_t) = \Theta_0 \Theta_0'$
 - Then we must have that

$$\Sigma_y - \frac{1}{\alpha^2} \text{Cov}(y_t, z_t) \text{Cov}(y_t, z_t)' \text{ is positive semi-definite} \quad (7)$$

Why? otherwise the leftover part is not a valid stochastic process (e.g., negative variance)

- The full argument simply applies the requirement (7) to all **frequencies** separately

High-frequency monetary shocks



High-frequency oil shocks

