LECTURE 1.3

VARs: Applications and Identification

ECON42240 Advanced Macroeconomics

Sam Deegan
University College Dublin

Spring 2027

Where We Are

Last Week's Lecture

1. **Specifying a VAR.**
   Which variables, how many lags, on what data.
2. **Estimating a VAR.**
   Equation by equation, by least squares.
3. **Reading the Output.**
   Impulse responses, error bands and variance shares.
4. **Residuals and Structural Shocks.**
   Why every mix of shocks fits the data.

Lecture Plan

1. **The Identification Problem.**
   Why residuals are not shocks.
2. **Recursive Identification.**
   What a Cholesky ordering assumes.
3. **Monetary Policy Shocks.**
   Why the ordering can decide the answer.
4. **Long-Run Restrictions.**
   Identifying technology shocks instead.

SECTION 1

The Identification Problem

A Residual Is Not a Structural Shock

Each residual (e_{it}) is a forecast error, not an economic event:

- **What Least Squares Returns:** The part of each variable the VAR's lags miss.
- **What Theory Defines:** A named disturbance, from policy, technology or preferences.
- **Why the Two Differ:** One economic surprise arrives in several equations at once.

> **A Structural Shock Has an Economic Name**
> A structural shock is an exogenous disturbance to one behavioural relation. Each residual mixes several such shocks.

Two Ways to Write the Same System

The structural form keeps the within-period effects (A_0) on the left:

$$A_0 Z_t = \sum_{i=1}^{p} A_i Z_{t-i} + \varepsilon_t \tag{1}$$

where:

- Z_t: vector of the n variables
- A_i: coefficients on lag i
- p: number of lags

Only the Reduced Form Can Be Estimated
Current variables on the right correlate with the equation's shock (ε_t).

What the Residuals Mix Together

Pre-multiplying by A_0^{-1} leaves only lags on the right, so least squares works:

$$Z_t = \sum_{i=1}^{p} A_0^{-1} A_i Z_{t-i} + \underbrace{A_0^{-1}\varepsilon_t}_{e_t} \tag{2}$$

Every Estimate Carries the Same Unknown Inverse

Least squares returns $A_0^{-1}A_i$ and $A_0^{-1}\varepsilon_t$, never A_i or ε_t. Each residual is a weighted sum of structural shocks.

One Lag Keeps the Algebra Short

With one lag, write the impact matrix (C) for A_0^{-1}:

$$Z_t = BZ_{t-1} + C\varepsilon_t \tag{3}$$

where:

- B: reduced-form lag matrix

Each Entry Is One Impact Effect

c_{ij} is shock j's effect on variable i in the period it arrives. Stacking more lags in companion form changes nothing here.

The Covariance Matrix Is the Only Evidence

Assume the structural shocks are uncorrelated, and scale each to unit variance:

$$\Sigma \equiv E(e_t e_t') = CC' \tag{4}$$

Estimation Stops at the Covariance Matrix

Least squares delivers the residual covariance (Σ) and nothing more about C. Identification is the set of assumptions supplying the rest.

Identification Needs $n(n-1)/2$ Restrictions

C holds more unknowns than the symmetric Σ holds distinct numbers:

$$\underbrace{n^2}_{\text{unknowns in } C} - \underbrace{n(n+1)/2}_{\text{numbers in } \Sigma} = n(n-1)/2 \qquad (5)$$

1

restriction with two variables

3

restrictions with three variables

Restrictions Come from Outside the Data

Each scheme closes the gap with a claim the sample cannot test:

- **Timing:** One variable cannot be changed within the period by another.
- **Sign:** A shock raises one variable and lowers another on impact.
- **The Long Run:** One shock has no permanent effect on one variable.

> **A Just-Identified Scheme Cannot Be Tested**
> Exactly $n(n-1)/2$ restrictions leave no slack. Every just-identified scheme fits the data equally well.

Naming a Shock Before Identifying It

The real oil price rose through the 2000s, and sharply again in 2022.

> **Question**
> What restriction would let you call the 2022 rise supply rather than demand? Name it.

Take answers from the room.

SECTION 2

Recursive Identification

An Ordering Supplies the Missing Restrictions

With two variables, a shock ordered second cannot enter the first residual:

$$\begin{pmatrix} e_{1t} \\ e_{2t} \end{pmatrix} = \begin{pmatrix} c_{11} & 0 \\ c_{21} & c_{22} \end{pmatrix} \begin{pmatrix} \varepsilon_{1t} \\ \varepsilon_{2t} \end{pmatrix} \tag{6}$$

Recursive Means Ordered by Within-Period Effects

A recursive system orders the variables so that A_0 is lower triangular. Its inverse, C, is then lower triangular too.

Cholesky Delivers the Triangular Matrix

An ordering fixes which entries of C are zero. One matrix fact fixes the rest:

$$\Sigma = PP' \tag{7}$$

where:

- P: lower-triangular Cholesky factor

> **One Factor, with Exactly the Missing Zeros**
> Only one such P has a positive diagonal, which pins down $C = P$. Its $n(n-1)/2$ zeros close the gap exactly.

Kilian Asks What an Oil Price Shock Is

Oil price run-ups precede many recessions, though oil is a small input. Kilian asks a prior question (Kilian, 2009):

- **The Old Question:** What are the effects of an oil price shock?
- **Kilian's Question:** What kind of oil price shock is this one?
- **Three Candidates:** Oil supply, global demand, and speculation in the oil market.

Kilian's Three-Variable Monthly VAR

Three monthly series, twenty-four lags, and an ordering argued from oil markets (Kilian, 2009):

$$z_t = (\Delta\mathrm{prod}_t,\ \mathrm{rea}_t,\ \mathrm{rpo}_t)', \qquad A_0 z_t = \sum_{i=1}^{24} A_i z_{t-i} + \varepsilon_t \tag{8}$$

where:

- z_t: Kilian's three variables
- $\Delta\mathrm{prod}_t$: growth of world oil output
- rea_t: global real activity index
- rpo_t: real price of oil

Kilian's Three Identifying Restrictions

Each zero is argued from how oil markets work; the data cannot test it:

$$A_0 = \begin{pmatrix} a & 0 & 0 \\ b & c & 0 \\ d & e & f \end{pmatrix} \tag{9}$$

1. **Oil Supply First:** Within the month, neither demand nor price changes oil production.
2. **Activity Second:** Oil production enters world demand at once, the price does not.
3. **Price Last:** Oil production and world demand both enter the price at once.

Counting Kilian's Identifying Restrictions

C holds nine unknowns and Σ six numbers, so three restrictions are missing:

> **Three Zeros Close a Gap of Three**
> Kilian's three zeros above the diagonal fill the gap exactly. No restriction is left over to test.

A different ordering therefore fits the same data equally well.

What the Residuals Mean Under This Ordering

A triangular A_0 makes C triangular. Each residual mixes its own shock and those ordered above it:

$$\begin{pmatrix} e_t^{\text{prod}} \\ e_t^{\text{rea}} \\ e_t^{\text{rpo}} \end{pmatrix} = \begin{pmatrix} c_{11} & 0 & 0 \\ c_{21} & c_{22} & 0 \\ c_{31} & c_{32} & c_{33} \end{pmatrix} \begin{pmatrix} \varepsilon_t^{\text{prod}} \\ \varepsilon_t^{\text{rea}} \\ \varepsilon_t^{\text{rpo}} \end{pmatrix} \tag{10}$$

Only the First Residual Is One Shock

The production residual is proportional to the oil supply shock. The other two residuals are the mixtures the ordering specifies.

Decomposing the Oil Price into Its Shocks

A historical decomposition splits each month's oil price among the identified shocks:

$$Z_t = C\varepsilon_t + BC\varepsilon_{t-1} + B^2C\varepsilon_{t-2} + \cdots \quad (11)$$

1. **One Shock at a Time:** Set the other two shock series to zero.
2. **Add the Three Paths:** The three paths and the initial conditions sum to the series.
3. **Read the Attribution:** Each month's price splits into three named contributions.

> **A Decomposition Needs the Realised Shocks**
> An impulse response traces one hypothetical shock. A decomposition feeds in every realised shock. Lecture 1.2 derived this moving-average form.

The Real Oil Price and Its Three Shocks

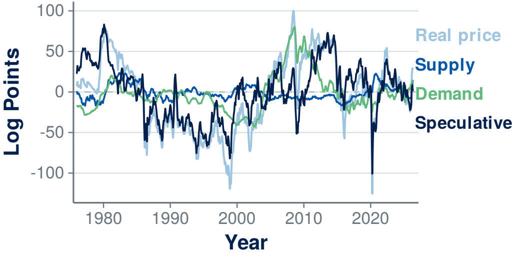

Figure 1: Cumulative effect of each identified shock on the real oil price, 1976–2026.

> **Supply Shocks Stay Small Throughout**
> Global demand accounts for most of the 2002–2008 climb. Speculative demand accounts for most of the 2022 spike.

What Kilian’s Three Shocks Explain

On his own sample, Kilian finds the three shocks matter very unequally (Kilian, 2009):

- **Supply Shocks:** Negligible for oil prices over time.
- **Speculative Shocks:** Can be significant for oil prices, limited for global activity.
- **The 2000s Run-Up:** Almost solely strong global demand.

> **An Oil Price Rise Has No Single Effect**
> An economy’s path after an oil price rise depends on its cause. The cause helps explain a 2000s oil climb with no recession.

Blanchard and Perotti on Fiscal Policy

Blanchard and Perotti close the same gap partly with outside evidence (O. Blanchard & Perotti, 2002):

- **Spending Is Slow:** Legislating a reply to a GDP surprise takes months.
- **Taxes Are Mechanical:** Estimated elasticities fix the same-quarter effect of GDP.
- **Output Is Last:** Both taxes and spending enter GDP within the same quarter.

GDP After a Tax Rise and a Spending Rise

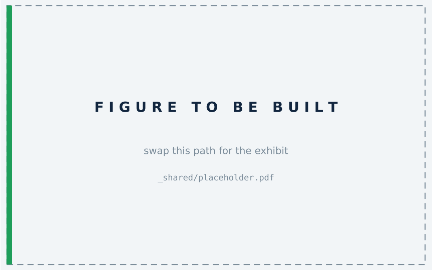

Figure 2: Blanchard and Perotti's GDP responses to a tax rise and a spending rise.

> **A Tax Rise Lowers GDP and a Spending Rise Raises It**
> The two paths differ in shape and in timing, not only in sign.

What Would a Monetary Policy Shock Be?

A central bank sets its rate after looking at the economy.

> **Question**
> Which part of a rate decision could count as a policy shock? Say why.

Take answers from the room.

SECTION 3

Monetary Policy Shocks

A Monetary Policy VAR Has Two Uses

Stock and Watson use a monetary VAR to answer two different questions (Stock & Watson, 2001):

- **The Scientific Question:** Only an exogenous policy change reveals policy's effects.
- **The Policy Question:** What follows a quarter-point rise beyond the bank's own rule?

> **Only the Part Outside the Rule Identifies an Effect**
> Much of a rate change is the bank following its own rule. Only the remainder identifies an effect.

Stock and Watson's Ordering

Inflation, unemployment and the funds rate, one restriction per row (Stock & Watson, 2001):

$$\begin{pmatrix} a_{11} & 0 & 0 \\ a_{21} & a_{22} & 0 \\ a_{31} & a_{32} & a_{33} \end{pmatrix} \begin{pmatrix} \pi_t \\ u_t \\ i_t \end{pmatrix} = A_1 Z_{t-1} + \cdots + A_p Z_{t-p} + \varepsilon_t \tag{12}$$

1. **Inflation First (π_t):** Sticky prices keep inflation free of within-quarter effects.
2. **Unemployment Second (u_t):** Current inflation enters unemployment, rates do not.
3. **The Funds Rate Last (i_t):** The Fed sees current inflation and unemployment.

Impulse Responses from the First Ordering

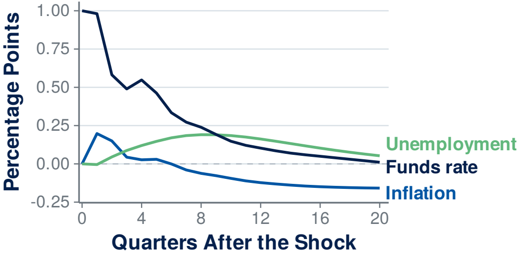

Figure 3: Responses to a one-point funds rate rise, first ordering, 1960–2007.

> **Inflation Rises for Five Quarters Before It Falls**
> Unemployment peaks about two years after the rise. Inflation stays above zero for five quarters, then falls below.

A Rate Rise Appears to Raise Inflation

In the first ordering, unemployment rises and inflation eventually falls, as theory predicts:

> **The Price Puzzle Is Inflation Rising First**
> The price puzzle is inflation rising for several quarters after a rate rise. Several VAR studies find the puzzle.

Whether the short-run rise is an effect or a failure is disputed.

Why the Price Puzzle Might Appear

The Fed watches more than a three-variable VAR contains, commodity prices especially:

- **What the Fed Sees:** Signals of future inflation that the small VAR omits.
- **What the VAR Records:** Rate rises arriving shortly before inflation rises.
- **The Confusion:** One omitted signal explains both, so correlation reads as causation.

> **Adding Commodity Prices Often Removes the Puzzle**
> Better specification, or data mining for the wanted result?

A Second Ordering, Equally Arguable

Three variables allow six orderings. Put the funds rate first and inflation last:

- **Why Rates First:** The Fed acts on lagged data, not on this quarter’s.
- **Why Inflation Last:** Prices can change at once when economic conditions change.

> **A Sensible-Looking Result Is Not Evidence**
> Which ordering a researcher keeps may depend on which results look sensible.

The Same VAR Under Two Orderings

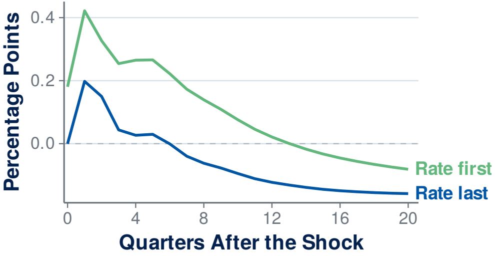

Figure 4: Inflation's response to a one-point funds rate rise, two orderings, 1960–2007.

Putting the Rate First Stretches the Puzzle
Data, lags and sample are identical. With the rate first, inflation stays above zero through quarter 13, not quarter 5.

Is This Data Mining?

Two orderings are defensible and one gives sensible answers. Adding commodity prices often removes the puzzle too.

> **Question**
> You have two orderings and one looks sensible. What would justify keeping the sensible one?

Take answers from the room.

Only Outside Evidence Can Rank Two Orderings

Both orderings fit the data identically, so the ranking needs outside evidence:

> **Report Both Orderings and Say Why**
> Report both sets of responses. State which ordering you prefer, and on what evidence.

Institutional timing, or an elasticity measured elsewhere, counts as evidence.

Rudebusch's Case Against Monetary VARs

Rudebusch argues that VAR monetary policy shocks are not policy surprises (Rudebusch, 1998):

- **The Rule Changed:** VARs impose one policy rule across decades of different regimes.
- **The Data Are Wrong:** VARs use revised figures the Fed never had.
- **The Information Is Thin:** The Fed watches far more than a VAR holds.
- **The Shocks Differ:** VAR shocks look nothing like fed funds futures surprises.

> **Similar Responses from Very Different Shocks**
> Models with quite different shock series report similar impulse responses. Rudebusch suspects data mining.

Sims's Reply to Rudebusch

Sims accepts the questions and rejects the conclusion (Sims, 1998):

- **Different Shocks:** Two models can size shocks differently and still agree on effects.
- **A Surprise Is Not Exogeneity:** A pre-announced rate rise is still exogenous policy.
- **The Faults Are General:** Linearity and variable choice afflict every macro model.

A Shock Measure Can Be Wrong and Still Useful

Sims argues an estimated effect can survive a mis-measured shock. The responses are what the exercise is for.

Romer and Romer Build Shocks from the Record

A narrative approach identifies policy shocks outside the VAR altogether (Romer & Romer, 2004):

1. **The Intended Rate:** Internal Fed documents give the policy rate the Fed intended.
2. **The Fed's Forecasts:** Internal forecasts of inflation and real activity, by meeting.
3. **The Residual:** Regress the intended rate change on those forecasts, keep the rest.

What the Narrative Shocks Missed

Coibion updates the series, and the recent recession raises a hard question (Coibion, 2012):

> **An Omitted Motive Looks Like a Shock**
> Large negative shocks mean cuts beyond what the forecasts justified. The Fed may have feared a financial crisis the regression omits.

Large stimulus before a slow recovery can then look ineffective.

Which Timing Restriction Would You Defend?

Every ordering so far rests on a claim about a single period.

> **Question**
> Of Kilian’s three timing restrictions, which would you defend? Which would you drop?

Take answers from the room.

SECTION 4

Long-Run Restrictions

When Timing Restrictions Are Guesswork

Recursive identification needs confident claims about one period. Often nobody has them:

- **Sometimes Defensible:** Oil output is physically slow; some data arrive with a lag.
- **Often Not:** Financial prices are set continuously, inside any period you pick.

> **Theory Predicts the Long Run, Not the Impact**
> Theory predicts that a demand shock leaves long-run output unchanged. Theory rarely predicts impact effects.

Restricting the Long Run Instead

Theory pins down where the economy ends up, not how it gets there:

- **A Positive Demand Shock:** Long-run output is unchanged, and the price level rises.
- **A Supply Shock:** A permanent productivity gain raises the long-run level of output.

> **A Long-Run Restriction Zeroes a Cumulative Effect**
> A long-run restriction sets one shock's total effect on one variable to zero. Impact effects are unrestricted.

Long-Run Effects Are Summed Responses

A level (y_t) changes by the running total of its growth rates (Δy_t):

$$y_{t+h} - y_{t-1} = \sum_{j=0}^{h} \Delta y_{t+j} \tag{13}$$

Summing Growth Responses Gives the Level Effect
A shock's responses are C, then BC, B^2C and so on. Their running sum is its effect on the level.

The Geometric Sum in Matrix Form

Summing every response is a geometric sum, and it has a closed form:

$$D = \left(I + B + B^2 + \cdots\right)C = (I - B)^{-1}C \tag{14}$$

where:

- D: long-run response matrix
- I: identity matrix

Each Entry Is a Cumulative Effect

d_{ij} is shock j's cumulative effect on variable i. The sum converges when every eigenvalue of B lies inside the unit circle.

The Blanchard–Quah Set-Up in Matrices

The long-run matrix is unknown, but DD' needs only least squares output:

$$DD' = (I - B)^{-1} \Sigma \left[(I - B)^{-1}\right]' \tag{15}$$

Only the Product Is Estimated

Estimates fix DD', as they fixed Σ. How DD' splits into D and its transpose is still open.

RESIDUAL COVARIANCE MATRIX

The Blanchard–Quah Long-Run Restriction

Two variables need one restriction, so assume D is lower triangular:

$$D = \begin{pmatrix} d_{11} & 0 \\ d_{21} & d_{22} \end{pmatrix} \tag{16}$$

- **The Zero:** The second shock has no long-run effect on the first variable.
- **What Stays Free:** The second shock may change the first variable for years.

Lower Triangular Again, in a Different Place

A recursive scheme makes A_0 lower triangular. A long-run scheme makes D lower triangular.

Uniqueness Delivers the Impact Matrix

A lower-triangular D is the unique Cholesky factor of DD', and C follows:

$$D = \text{chol}\left((I - B)^{-1}\Sigma\left[(I - B)^{-1}\right]'\right), \qquad C = (I - B)D \tag{17}$$

The Long Run Pins Down the Impact Effects

Restricting cumulative effects delivers C, so every impulse response follows. No timing assumption was needed.

Supply and Demand Shocks, Identified

Blanchard and Quah applied the method to output growth and unemployment (O. J. Blanchard & Quah, 1989):

- **The Variables:** The log-difference of GDP, and the unemployment rate in levels.
- **The Labels:** The shock with a permanent output effect is supply, the other demand.
- **Their Finding:** Demand shocks account for most short-run output fluctuation.

> **The Labels Are Assumed, Not Discovered**
> Nothing in the data makes the permanent shock a supply shock. The restriction assigns the name.

Supply and Demand Shocks on 1948–2019 Data

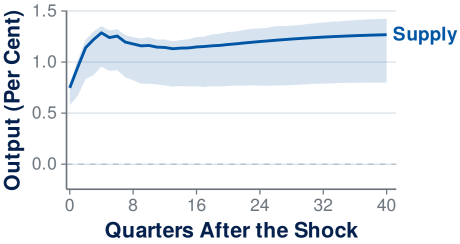

Figure 5: Supply shock: cumulative output response, 1948–2019.

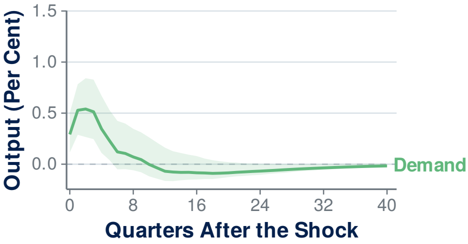

Figure 6: Demand shock: cumulative output response, 1948–2019.

> **Demand Carries the Short Run, Supply the Level**
> The demand shock's long-run output effect is zero by construction. The supply shock's settles near 1.3 per cent.

Galí's Technology and Hours VAR

Galí orders productivity growth first, so only technology changes long-run productivity (Galí, 1999):

$$Z_t = (\Delta z_t,\ \Delta n_t)', \qquad D = \begin{pmatrix} d_{11} & 0 \\ d_{21} & d_{22} \end{pmatrix} \tag{18}$$

where:

- Δz_t: growth of output per hour
- Δn_t: growth of hours worked

> **A Technology Shock Changes What Labour Produces**
> Only a technology shock permanently changes output per hour. Both shocks' long-run effects on hours are free.

Why Galí's VAR Tests the RBC Model

Real business cycle models make technology shocks the source of the cycle:

> **The RBC Model Predicts Hours Rise with Technology**
> Working when you are productive beats working when you are not. A positive technology shock raises hours and output.

An identified VAR delivers the sign of the hours response directly.

Responses to Technology and Non-Technology Shocks

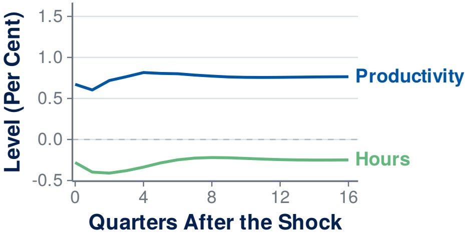

Figure 7: Technology shock: productivity and hours, 1948–2019.

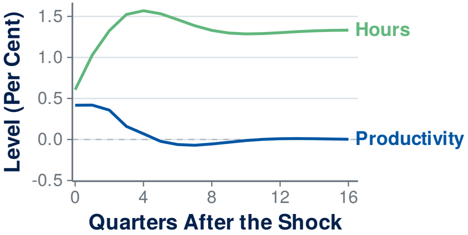

Figure 8: Non-technology shock: productivity and hours, 1948–2019.

> **Hours Fall After a Positive Technology Shock**
> Productivity rises permanently, and hours are still lower four years on. The RBC model predicts hours rise.

Reading Galí's Two Impulse Responses

Neither shock produces what a technology-driven cycle predicts (Galí, 1999):

Demand-Determined Output Explains the Falling Hours

More efficiency means demanded output needs less labour. Short-run output looks demand-driven.

Non-technology shocks raise output and productivity together in the short run.

What Galí's Results Cost the RBC Model

Galí's own interpretation runs against the real business cycle account (Galí, 1999):

- **Labour Adjustment Costs:** In a boom, work existing staff harder rather than hire.
- **In a Recession:** Employed labour is under-utilised, so measured productivity falls.

> **The Hours Finding Is Contested**
> Later work specifies hours differently and gets different answers. The question is not settled.

What Every Identification Scheme Shares

Orderings, outside elasticities, narrative records and long-run restrictions all do one job:

- **The Common Job:** Each supplies from outside the sample what the residuals cannot.
- **The Common Cost:** Each restriction is an assumption the sample cannot test.

Report the Scheme, Not Only the Responses
An impulse response without its identifying assumption is not a finding. The assumption is the claim.

Before Lecture 1.4

1. **Play:**
   Re-identify last week's VAR under two orderings and compare the responses.
2. **Apply:**
   Specify and identify one VAR, then save it for later weeks.
3. **Read:**
   Whelan on latent variables, then Kilian's oil paper in full (Kilian, 2009).

Fix the VAR's series, sample, lag length, ordering and file, and keep all five.

Next Lecture

Lecture 1.4 turns to variables no dataset contains, and to the Kalman filter:

1. **Unobserved Components.**
   Writing a model in state-space form.
2. **The Kalman Filter.**
   Extracting a state you cannot see.
3. **Estimation and Smoothing.**
   Estimating the variances, and revising the past.
4. **The Natural Rate of Interest.**
   What the filter is used for.

Thank you

sam.deegan@ucdconnect.ie
sam-deegan.com

BACKUP

Appendix

Notation

Symbol	What it is	First met
e_{it}	residual	Lecture 1.1
A_0	within-period effects	Lecture 1.2
Z_t	vector of the n variables	Two Ways to Write the Same System
A_i	coefficients on lag i	Two Ways to Write the Same System
p	number of lags	Two Ways to Write the Same System
n	number of variables	Lecture 1.2
ε_t	structural shock	Lecture 1.2
C	impact matrix	Lecture 1.2
B	reduced-form lag matrix	One Lag Keeps the Algebra Short
Σ	residual covariance	Lecture 1.2
P	lower-triangular Cholesky factor	Cholesky Delivers the Triangular Matrix
z_t	Kilian's three variables	Kilian's Three-Variable Monthly VAR

Notation

Symbol	What it is	First met
Δprod_t	growth of world oil output	Kilian's Three-Variable Monthly VAR
rea_t	global real activity index	Kilian's Three-Variable Monthly VAR
rpo_t	real price of oil	Kilian's Three-Variable Monthly VAR
π_t	inflation	Stock and Watson's Ordering
u_t	unemployment	Stock and Watson's Ordering
i_t	funds rate	Stock and Watson's Ordering
y_t	level	Long-Run Effects Are Summed Responses
Δy_t	growth rates	Long-Run Effects Are Summed Responses
D	long-run response matrix	The Geometric Sum in Matrix Form
I	identity matrix	The Geometric Sum in Matrix Form
Δz_t	growth of output per hour	Galí's Technology and Hours VAR
Δn_t	growth of hours worked	Galí's Technology and Hours VAR

References

Blanchard, O. J., & Quah, D. (1989). The dynamic effects of aggregate demand and supply disturbances. *American Economic Review*, *79*(4), 655–673.

Blanchard, O., & Perotti, R. (2002). An empirical characterization of the dynamic effects of changes in government spending and taxes on output. *The Quarterly Journal of Economics*, *117*(4), 1329–1368.

Coibion, O. (2012). Are the effects of monetary policy shocks big or small? *American Economic Journal: Macroeconomics*, *4*(2), 1–32.

Galí, J. (1999). Technology, employment, and the business cycle: Do technology shocks explain aggregate fluctuations? *American Economic Review*, *89*(1), 249–271.

Kilian, L. (2009). Not all oil price shocks are alike: Disentangling demand and supply shocks in the crude oil market. *American Economic Review*, *99*(3), 1053–1069.

Romer, C. D., & Romer, D. H. (2004). A new measure of monetary shocks: Derivation and implications. *American Economic Review*, *94*(4), 1055–1084.

Rudebusch, G. D. (1998). Do measures of monetary policy in a VAR make sense? *International Economic Review*, *39*(4), 907–931.

References

Sims, C. A. (1998). Comment on Glenn Rudebusch's "Do measures of monetary policy in a VAR make sense?" *International Economic Review*, *39*(4), 933–941.

Stock, J. H., & Watson, M. W. (2001). Vector autoregressions. *Journal of Economic Perspectives*, *15*(4), 101–115.