Impulse responses

Christopher A. Sims
Princeton University
sims@princeton.edu

March 1, 2020

Interpreting them

With a VAR model written in stacked form as

$$(I - BL)y_t = \varepsilon_t \ ,$$

with ε_t independent of or uncorrelated with $\{y_{t-s}, s > 0\}$, we can represent y_t as

$$y_{t+s} = \sum_{v=0}^{s-1} B^v \varepsilon_{t+s-v} + B^s y_t$$

whether or not y_t is stationary (i.e. whether or not all the eigenvalues of B are less than one in absolute value.)

Interpreting them

We'll now focus on the case where y is stationary, so $B^s y_t \to 0$ as $s \to \infty$, and the model therefore has the moving average representation

$$y_t = C(L)\varepsilon_t \,.$$

Almost everything we say about this case applies also in the non-stationary case.

The i, j element of the C_s matrix can be interpreted as the change in the forecast of $y_{i,t+s}$ induced byh a unit change in the j'th element of ε. This makes C_{ijs}, treated as a function of s, an "impulse response function".

Orthogonalization

- Generally $\mathrm{Var}(\varepsilon_t) = \Sigma$ is not diagonal, so the variable innovations are correlated.

- Impulse responses are easier to interpret if they are responses to shocks that are uncorrelated.

- There are many ways to tranform ε_t so it has a diagonal, or even identity, covariance matrix. The most widely applied one in descriptive time series models is based on the Cholesky decomposition of Σ: $W'W = \Sigma$ with W upper triangular.

- Then we can write $\varepsilon_t = W'\zeta_t$, where $\mathrm{Var}(\zeta_t) = I$ and

$$y_t = C(L)W'\zeta_t\,.$$

The i, j element of $C_s W$ is now the response of variable i to the j'th orthogonalized shock ζ_j.

Interpreting the Cholesky-orthogonalized shocks

- ζ_1 is just ε_{1t} divided by its standard deviation.
- ζ_2 is the residual in a regression of ε_{2t} on ε_{1t}, divided by its standard deviation.
- etc.
- If one suspects a causal ordering, with one variable moving first, then moving the next, etc., one can reflect that guess by ordering the variables earlier in the ordering first.

Variance decompositions

- With the residuals orthgonalized, one can split the variance of each variable y_{it} into components that add up to the total variance, with each component generated by one of the ζ's. (This can't be done in the non-stationary case.) The component of the variance in y_{it} attributable to shock j is $\sum_{s=0}^{\infty} G_{ijs}^2$, where $G(L) = C(L)W'$.

- Even in the non-stationary case, we can break up the variance in the k-step ahead forecast error in y_{it}, which is just $\sum_{s=0,j=1}^{s=k-1,j=n} G_{ijs}^2$, into components attributable to the ζ's.