Error Bands for Impulse Responses

Christopher A. Sims
Princeton University
sims@princeton.edu

November 14, 2019

Error bands for impulse responses

- In one sense, this is straightforward: make draws from the posterior pdf of Σ and A (orβ) and for each draw calculate impulse responses $c_{ij}(t)$
- Plot the $c_{ij}t$ corresponding to the MLE (and/or the mean or median of the draws), together with the 5% upper and lower tails of the draws, for example. Could also plot HPD intervals.

These are not confidence intervals

- It's a Bayesian calculation. It gives intervals with a clear interpretation, by a straightforward procedure, but they are not confidence intervals (except asymptotically!).

- True confidence intervals for individual c_{ij}'s are not possible.

- This is a special case of a general point. If θ is the complete parameter vector for the distribution of the data X, then we can *always* (in principle) produce a 90% (say) confidence *region* for θ by constructiong a 90% significance level test for θ as H_0 for each θ in the parameter space. The set of θ's that are accepted in a given sample is an exact 90% confidence set.

- But if we try to construct a confidence set for an individual θ_i, the problem is that the distribution of any test statistic generally depends on all the parameters, not just θ_i.

- The normal linear regression model is a special case where there is a set of test statistics that depend on single θ_i's.

So what are frequentist CI's?

- There are "asymptotic" confidence intervals. These have coverage probabilities that converge to the correct ones for parameter values in some neighborhood of the true parameter value. They can be constructed by linearizing the mapping from B_{ij} (the AR coefficients) to C_{ij}, then transforming the normal asymptotic distribution for B to the corresponding approximate normal distribution for C.

- But these have no more frequentist asymptotic justification than do the Bayesian intervals.

- As the forecast horizon expands, the nonlinearities rapidly become more extreme, so the intervals based on linearization are always inaccurate at distant horizons.

- The Bayesian intervals are accurate whether or not A might have roots of one or larger in absolute value. The frequentist intervals are uninterpretable if that is true.

Bootstrap bands

- Bootstrap bands. Because of the well-known bias of MLE estimates, the simple parametric bootstrap gives terrible results. So a "double bootstrap" is needed. The first estimates the bias in $\hat{\rho}$, the second bootstraps again, using a bias-corrected estimate of ρ to generate the draws.

- The bootstrap bands generated this way are invalid if any roots exceed or equal 1, but draws will sometimes occur in that region.

- Such draws have to be discarded, but if there are a lot of them (as is likely when the VAR is large) this distorts results.

Delta-method bands

- With the VAR put into stacked first-order form

$$y_t = c + Ay_{t-1} + \varepsilon_t \,,$$

  the irf's are elements of A^s, $s = 0, \ldots, \infty$.

- With y stationary $\hat{A}$ by OLS is asymptotically normal.

- $$\frac{\partial A^2}{\partial A} = (I \otimes A)\vec{A} + (A' \otimes I)\vec{A} \,.$$

- Explicit formulas for higher powers of A are also available, allowing calculation of delta-method asymptotics.

- But all roots must be below one in absolute value, and if the are, the error bands shrink exponentially with horizon, often very fast.

Pointwise bands, alternatives

- These bands are constructed from marginals for each c_{ij}.
- They cannot answer a question like "How likely is a hump-shaped response?"
- One could directly get a distribution for "up-life" from the MCMC sample.
- SZ alternative: display $1 - \alpha$ probability bands for principal components of $\mathrm{Var}(c_{ij\cdot})$.
- Like displaying bands for principle components of the posterior of β in a regression, but with a more direct intuitive interpretation of the bounds.

Do we want bands that include the full response?

- Consider a multivariate regression. Would you rather have t-stats for each coefficient, or a set of intervals, one for each variable, such that *all* the coefficients are inside the corresponding multivariate rectangle.

- The latter gets to be a very big set if many coefficients are included, much bigger than that defined by the usual coefficient-by-coefficient confidence intervals.

Geometry break

Joint bands for pairs of impulse responses via the PC approach

- Using the PC approach, we can ask, e.g., whether large upward deviations in the GDP response to a monetary policy shock are associated with large negative deviations in the price response.

- One does this by stacking the two responses into a longer vector, calculating its covariance matrix from the MCMC draws, and extracting eigenvectors and eigenvalues.