Linearizing and solving multivariate RE models

Christopher A. Sims
Princeton University
sims@princeton.edu

April 5, 2020

Road map

- We are aiming to be able to discuss the models and results in the papers listed in Topic 4, *FTPL in Keynesian Models*, in the newly revised version of the course syllabus.
- In the models we have looked at so far, it has been possible to reduce the main work of solution to studying a one-dimensional differential equation. In other words, these models could be arrange to have a single state variable.
- Realistic New Keynesian models with active fiscal policy, in which raising interest rates reduces inflation, the price level has smooth time paths, and increasing primary deficits increases real debt, require more complex models.

The canonical continuous time, linear, perfect foresight model

$$\Gamma_0 Dy = \Gamma_1 y + C$$
$$\Theta\Gamma_0 y_0 = z_0 \ .$$

The Θ matrix is usually a selection matrix (all zeros but for a single 1 in each row) that defines the backward-looking equations. The z_0 vector is the vector of values inherited from the past at time zero. D is the differential operator: $Dy = \dot{y}$, $D^2 y = \ddot{y}$, etc.

What `gensys` will do for you

- You give it a system in the canonical form.
- It returns a system in the form

  $$Dy_t = Gy_t + c$$

  whose solution paths for y that satisfy the original equation system, but also have the property that their solution paths are stable.
- It does this by enforcing relations among the elements of the y vector.
- One can think of this as pinning down the rest of the y_0 vector, the part left indeterminate by the $\Theta\Gamma_0 y_0 = z_0$ condition.

Additional details

- There may be no solution (non-existence) of this form, and it may also be that enforcing stability does not generate enough linear restrictions to fully determine y_0 (non-uniqueness). `gensys` will flag these situations.

- If (as might be true if your stability condition comes from transversality) you want to allow exponential growth, so long as it is below some rate β, `gensys` will let you specify β.

Preprocessing

- If you are starting with a nonlinear model, finding the steady state and calclulating the derivatives needed for a linearization can be a lot of work, prone to error.

- The `gensys` package includes programs that will let you specify a model as a set of nonlinear equations, labeled if forward-looking, then let the computer find the steady state and construct the linearization.

- The DYNARE system will do the same thing, and more, for discrete time models.

- DYNARE works in Matlab. The current `gensys` works in R.

How is it done?

We'll show how it works in the special case where Γ_0 is non-singular. There can be perfectly reasonable systems with singular Γ_0, and `gensys` handles them properly, but that more general case makes for messy algebra.

Multiplying the system by Γ_0^{-1} leads us to

$$Dy_t = \Gamma_0^{-1}\Gamma_1 y_t + \Gamma_0^{-1}C = Ay_t + h\,.$$

We'll avoid some more algebra by assuming that A has a Jordan decomposition

$$A = P\Lambda P^{-1}$$

with Λ diagonal. In general, there can be ones on the first diagonal above the main diagonal in the Jordan decomposition, and the software handles that situation.

Transforming to a stack of univariate equations

Multiplying the system on the previous slide on the left by P^{-1} and letting $x = P^{-1}y$ leads us to

$$Dx_t = \Lambda x_t + P^{-1}h\ .$$

This is now a stack of univariate linear ODE's (ordinary differential equations). If the i'th element of the diagonal of Λ is λ_i, then those equations in which $\lambda_i > 0$ imply time paths for $x_i(t)$ that explode exponentially unless $x_i(t) \equiv -P^{i\cdot}h$, where $P^{i\cdot}$ is the i'th row of P^{-1}.

Deriving stability conditions

Let's arrange the equations so that all the $\lambda_i > 0$ cases are at the bottom. We'll use the notation $P^{2\cdot}$ for the rows of P^{-1} that correspond to these unstable roots $\lambda_i > 0$. Then to enforce stability on the solution, we require

$$P^{2\cdot} y \equiv -P^{2\cdot} h \,.$$

This, along with the initial conditions, let us fully determine y_0 via

$$\begin{bmatrix} \Theta\Gamma_0 \\ P^{2\cdot} \end{bmatrix} y_0 = \begin{bmatrix} z_0 \\ -P^{2\cdot} h \end{bmatrix} \,.$$

Existence and uniqueness

The last equation on the previous slide, which was said to "determine y_0", obviously can only do so if the matrix

$$\begin{bmatrix} \Theta\Gamma_0 \\ P^{2\cdot} \end{bmatrix}$$

is square and non-singular. In particular, the matrix must have full column rank if it is to fully determine y_0, and full row rank if it is to yield a solution for y_0 for arbitrary z_0 and h. That the system determines y_0 uniquely is the "uniqueness" property of the solution, and that it has a solution at all for arbitrary initial conditions is the "existence" property.

Root counting

In most models, we fail to have existence if the number of "excessively explosive" roots is more than n, the total number of equations, minus m, the number of backward-looking equations. We fail to have uniqueness if the number of excessively explosive roots is less than $n - m$. But these are just rules of thumb to ensure the matrix is square. One can't be sure about this without calculating the matrix itself to check rank.

Exogenous shocks

When we consider discrete time systems, we will include stochastic disturbances from the start, even in nonlinear systems. For continuous time nonlinear systems, adding stochastic disturbances takes us into technical weeds that we have avoided. But for linear systems, adding white noise shocks is pretty easy.

What is white noise?

- It would be a serially uncorrelated process ε_t, such that $E[\varepsilon_t \varepsilon_s] = 0$ for all $t \neq s$, except that there is no such thing.

- We use a white noise process ε_t to define ordinary processes that have values at given times by averaging ε_t over intervals. So a standard white noise process is an ε_t satisfying

$$E\left[\int_a^b \varepsilon_t \, f d dt \cdot \int_c^d \varepsilon_t \, dt\right] = \text{length}([a,b] \cap [c,d]) \, .$$

for any pair of intervals.

ODE's with white noise disturbances

If we tack a white noise onto the end of a univariate first order ODE, we get

$$\dot{y} = -ay + \varepsilon_t \;, \text{ whose stable solution is, if } a > 0$$

$$y_t = \int_0^\infty e^{-as} \varepsilon_{t-s} \;.$$

The y_t defined this way is an ordinary stochastic process with continuous (but not differentiable) time paths. It satisfies $\mathrm{Cov}(y_t, y_{t-s}) = (1/a^2)e^{-|s|}$.

Adding ε's to our perfect-foresight models

- We can add white noise disturbances to the right-hand sides of the equations defining our linear rational expectations systems, and interpret our solutions as showing how y is "driven" by the exogenous disturbances.

- The solution mechanics are unaffected by whether or not we include disturbances.