

Solution methods for DSGE's

April 13, 2009

The general framework

$$\underset{n \times 1}{K} \left(\underset{n \times 1}{w_t}, \underset{n \times 1}{w_{t-1}}, \underset{m \times 1}{\sigma \varepsilon_t} \right) + \underset{p \times 1}{\Pi} \sigma \eta_t = 0,$$

or, equivalently

$$Q_1 K(w_t, w_{t-1}, \sigma \varepsilon_t) = 0$$

$$E_t[Q_2 K(w_{t+1}, w_t, \sigma \varepsilon_{t+1})] = 0$$

$$\begin{bmatrix} Q_1 \\ Q_2 \end{bmatrix} \text{ full rank, } Q_1 \Pi = 0, \quad |Q_2 \Pi| \neq 0.$$

What is a solution?

One definition: A function $F(w_{t-1}, \sigma \varepsilon_t)$ such that when $w_t = F(w_{t-1}, \varepsilon_t)$, the equations of the model, plus any side constraints (transversality, feasibility), are satisfied.

Sometimes we split the vector w , so $w_t = (s_t, c_t)$ and the solution is written as

$$s_t = \hat{K}(s_{t-1}, c_{t-1}, \varepsilon_t)$$
$$c_t = \hat{F}(s_t).$$

In this case s is called the “state” and, if the problem is a single-agent optimization, c is called the “control”. In economics in the single agent case, F is often called the “decision rule”.

General points

- All numerical methods are approximate. There is no automatic presumption that “nonlinear” methods will be more accurate than even linearization, which is one approximate method.
- There are local (based on Taylor expansions) and global methods. There is again no automatic presumption that global methods are more accurate — even “globally more accurate” than local ones.
- Choosing a method for measuring the accuracy of solutions is at least as important as choosing a solution method, since no choice of method guarantees accuracy.

- Are we interested in decision rules and pricing functions directly, or in simulated time paths for variables in the model?

Accuracy measures

Given F , check Euler equation errors. That is, check whether

$$E_{t-1}K(F(w_{t-1}, \varepsilon_t), w_{t-1}, \varepsilon_t) = 0 .$$

Two ways to check:

- At a priori fixed grid of w values
- At grid weighted by frequency of visits to state

Fixed grid

At each grid point w , calculate

$$\int K(F(w, \varepsilon), w, \varepsilon)p(\varepsilon)d\varepsilon$$

This should be zero. It helps to have the solution in a form that provides $c = \hat{F}(s)$, since then our grid can be over the lower dimensional space where s resides, rather than the whole w space.

What is a small error? When comparing two models, this is not an issue, but when deciding whether the solution to a given model is “accurate”, this is critical.

Sometimes the expectational errors have the character of a mistaken

forecast of an expected rate of return, which provides a natural metric for measuring accuracy.

How to pick the grid? Without simulating the model, it is possible to be quite mistaken about which values of w or s occur often. Accuracy at values of w that are very rare may be much less important than accuracy in commonly visited regions. On the other hand, when considering interventions that change the nature of the w process, we may need accuracy in otherwise seldom visited regions of w -space.

Checking accuracy by simulation

One can start up the model from some reasonable initial condition (near deterministic steady-state, e.g.), run it for N periods, generating a $\{w_t\}$ sequence. Check accuracy by calculating

$$\frac{1}{N} \sum_1^N K(w_t, w_{t-1}, \varepsilon_t)$$
$$\frac{1}{N} \sum_1^N g(w_{t-1}) K(w_t, w_{t-1}, \varepsilon_t), \text{ etc.,}$$

where g is an arbitrary function of w_{t-1} . These should all converge to zero at an exact solution. Or more systematically, run a regression of the realized time series of $K(w_t, w_{t-1}, \varepsilon_t)$ values on a constant and a collection $\{g_i\}$ of functions of lagged (possibly lagged more than once) w values.

Comparing two models, it is fairly straightforward to see which is better. But an absolute criterion for when a solution is “accurate” is difficult to formulate. With the regression check, one might say that if the simulated samples (from many initial w 's) are of a length similar to historical time series data series, accuracy is pretty good if an F -test for the null hypothesis that all coefficients are zero is passed at the 5% or 1% level. This suggests that whatever errors agents are making, they would have a hard time themselves in detecting the errors. (Assuming they are no cleverer than we are in selecting $\{g_i\}$.)

Choice of accuracy measure

- Checking by simulation tends to weight accuracy at different values of the state by how frequently the values occur. This is usually an advantage.
- If we plan to use the solution to examine changes in parameters that change the ergodic distribution of the states, accuracy assessed by simulation at the original parameter values may be misleading.
- Assessing accuracy by simulation amounts to doing numerical integration by Monte Carlo methods. Often (Judd would say always) there are numerically more efficient methods.

Backsolving

1. Guess a \hat{F} for $c_t = \hat{F}(s_t)$, and generate an η_t process that satisfies $E_t \eta_{t+1} = 0$.
2. Simulate a system formed from

$$Q_2 K(w_t, w_{t-1}, \varepsilon_t) = \eta_t$$

$$c_t = \hat{F}(w_t)$$

$$M s_t = M \hat{K}(w_{t-1}, \varepsilon_t),$$

where M selects some equations that make the equation count of the full system match the length of w .

Some subvector of ε_t , maybe even the whole vector, is solved for rather than independently generated. The simplest case is where the third block

of equations is not needed to solve for w_t and ε_t does not appear in the first block. Then the first two blocks can be used to solve for the w_t path, and the state equation block is then used to solve for the ε_t path.

Limitations of backsolving

- This may or may not produce a solution that satisfies boundary conditions. (Lack of explosiveness, no-Ponzi, etc.) If it does, the solution certainly satisfies the FOC's and constraints of the original system — except that the parts of the stochastic process ε_t not included in the system now have properties different from those originally postulated — they may be serially correlated, have non-zero means, etc.
- If non-linearities are not too strong, the distortion of the ε properties may not be very great.
- In any case, the properties of ε_t are usually specified ad hoc — economic theory usually has little to say about the properties of ε .

- The main limitation here is that the “guess” of \hat{F} in many cases must be quite good in order for the system to be stable. One possibility: Obtain a linear \hat{F} by log-linearizing the system; use that in the backsolved nonlinear system. This is likely to produce a stable solution if the shocks are not too large or nonlinearity is not too severe.

Local methods

- Form a Taylor expansion of K in the neighborhood of deterministic steady state, of order n .
- Solve the resulting system to n 'th order.
- That is, find an n 'th order polynomial F that sets the first n terms of the Taylor expansion of $E_{t-1}K$ to zero.
- Such solutions can be locally accurate, in the sense that as initial deviations of the state from steady state and the size of shocks shrink, the relative accuracy of the solution increases.

Limitations of local methods

- They won't work unless K is differentiable to order n and shocks have n finite moments. (The conditions are actually stronger than just the existence of the moments.)
- Though the calculations are closed form, so there is no iteration for a given n , they can involve staggering amounts of algebra. Note that the fourth derivative tensor for a system with four variables contains $4^4 = 256$ elements and the set of all derivatives up to fourth order has 340. Computing each of these by hand without error may be effectively impossible. High order expansions beyond second order are therefore usually practical only with the aid of computer algebra or “automatic differentiation”.

- For systems with, say, 10 variables and several lags, storage space in the computer can become an issue, and in such systems efficient computation requires use of sparse matrix methods.
- One iteration, but programming that one iteration may take longer than thousands of iterations of another approach.
- If K is sufficiently differentiable, the solution may become perfectly accurate as n increases without bound, but only within some “radius of convergence”, and it is difficult to know in advance what this radius is.

Parametrizing something, then optimizing or equation-solving

— 4/14 lecture —