EC316a: Advanced Scientific Computation, Fall 2003

Notes Section 4

Discrete time, continuous state dynamic models: solution methods

We consider now solution methods for discrete–time models in which decisions are made over continuous state variables. These models give rise to functional equations whose unknowns are entire functions defined on a subset of Euclidean space: for instance, the unknown of the Bellman equation

$$V(s) = \max_{x \in X(s)} \{f(s,x) + \delta E V(g(s,x,\epsilon))\}$$

is the value function $V(\cdot)$. In most applications, these functional equations lack closed–form solutions, and can only be solved by numerical approximation methods. One numerical method widely used by economists is linear–quadratic–[Gaussian] (LQG) approximation, but it is often inadequate to deal with the sorts of models which may be encountered. A more generally useful method, developed in the physical sciences, is that of collocation, which is flexible, accurate, and numerically efficient. We will now describe the collocation method, after presenting the LQG method.

Linear–quadratic–[Gaussian] control

The linear–quadratic control model is an unconstrained Markov decision model with a quadratic reward function $f(s, x)$ and a linear state transition function $g(s, x, \epsilon)$. It is of special importance because it is one of the few continuous state Markov decision models known to have a finite–dimensional solution: the optimal policy and shadow price functions of the infinite–horizon LQ control problem are both linear in the state variable. The parameters of the shadow price function are characterized by nonlinear fixed–point Riccati equations, which may be solved for each period with backward recursion. Thus standard nonlinear equation solution methods may be used to solve a LQ problem.

The LQ model may also be applied in a stochastic context, with "well-behaved" (Gaussian) errors, since the shadow price and optimal policy functions depend only on the mean (or expected value) of the shock to each state, rather than on the variance or higher moments of the shock process. This gives rise to "certainty equivalence", in which one may replace the expectation of the shock with its mean, and solving the resulting deterministic problem. The problem is then known as an LQG problem.

The LQ methodology is often applied to more complex models (e.g. those not possessing a linear state transition function: for instance, a macroeconomic model that cannot be linearized) by replacing the nonlinear f and g functions with linear and quadratic approximants, and solving the resulting LQ problem. These approximations are characteristically derived by forming first– and second–order Taylor expansions around the certainty–equivalent steady state of the model. However, that methodology requires that the steady state of the model is computable, and that Taylor expansions around that steady state will not experience too much curvature in the relevant functions. Likewise, this linearization process usually discards any constraints on the states and actions. If those constraints are binding in the neighborhood of the steady state (e.g., the Federal Reserve's ability to lower interest rates at this point in time is limited), then the LQ approximation will be particularly poor. Therefore, methods which do not rely on these Taylor approximations are to be preferred.

Bellman equation collocation methods

Consider an infinite–horizon discrete–time model with a one–dimensional state and action space and univariate shocks. In a continuous state context, this model has the Bellman equation

$$V(s) = \max_{x \in X(s)} \{f(s,x) + \delta E V(g(s,x,\epsilon))\}$$

Assume that the state space is a bounded interval S of the real line, and that the actions are either discrete or continuous and subject to simple bounds $a(s) \leq x \leq b(s)$ that are continuous functions of the state.

To compute an approximate solution using collocation methods, we first write the value function approximant as a linear combination of n known basis functions $\phi_1, \phi_2, \ldots, \phi_n$ on S with undetermined coefficients:

$$V(s) \approx \sum_{j=1}^{n} c_j \phi_j(s)$$

Second, fix the basis function coefficients $c_1, c_2, \ldots, c_n$ by requiring the value function to satisfy the Bellman equation at n collocation nodes $s_1, s_2, \ldots, s_n$. This strategy replaces the Bellman functional equation with a system of n nonlinear equations in n unknowns:

$$\sum_{j=1}^{n} c_j \phi_j(s_i) =$$

$$\max_{x \in X(s_i)} \{f(s_i, x) + \delta E \sum_{j=1}^{n} c_j \phi_j(g(s_i, x, \epsilon))\}$$

which may be compactly expressed in vector form as the collocation equation $\Phi_C = v(c)$ where Φ_C, the collocation matrix, is an n x n matrix whose typical i j element is the j^{th} basis function evaluated at the i^{th} collocation node. When v, the collocation function, is evaluated at a particular vector of basis coefficients c, it yields a vector whose i^{th} entry is the value obtained by solving the optimization problem embedded in the Bellman equation at the i^{th}

collocation node s_i, replacing the value function V with its approximant $\sum_j c_j\phi_j$.

In principle, the collocation equation may be solved with any nonlinear equation solution method: as a fixed–point problem $c = \Phi^{-1}v(c)$, or as a root–finding problem $\Phi_C - v(c) = 0$, solving for c using Newton's method or a quasi–Newton method.

If the problem is stochastic, expectations must be computed in a practical manner: for instance, by replacing the continuous random variable ϵ in the state transition function g with a discrete approximant: for instance, one that assumes values $\epsilon_1, \epsilon_2, \ldots, \epsilon_k$ with probabilities $w_1, w_2, \ldots, w_k$ respectively. This is known as a numerical quadrature method, which computes discrete integrals as approximations to the continuous probability density function.

A number of practical decisions must be made when applying the collocation method. The basis functions (interpolators, such as spline functions) and collocation nodes must be selected, and an algorithm for solving the collocation equation chosen. Last, an appropriate quadrature technique for dealing with expectations must be selected. The choice of basis–node scheme will depend on the curvature of the value function. The larger the number of basis functions and collocation nodes, the greater the computational burden, so the researcher will want to experiment with various basis–node schemes and dimensions of the problem to render it computationally efficient.

Collocation methods address many of the shortcomings of LQ approximations, since they employ global (rather than local) function approximation schemes, and are not limited to the first– and second–degree approximations afforded by LQ methods. However, polynomial and spline approximants used in collocation methods can behave strangely, especially in the presence of nondifferentiabilities in the value function and binding constraints on the action variables (which may cause problems of nonconvergence).

Although we have discussed collocation methods for a very simple problem, the routines included in the CEtools MATLAB toolbox will support solution of models with multidimensional states, actions, and shocks. The major computational challenge in implementing these methods for any model is the implementation of the `vmax(s,c)` function: a function that solves the optimization problem embedded in the Bellman equation at the collocation nodes and returns the collocation function values and derivatives. The `vmax` function will return several objects: an n x 1 vector `v` of optimal values at the collocation nodes, an n x 1 vector `x` of associated optimal actions at the nodes, and an n x n matrix `vjac`, the Jacobian of the collocation function evaluated at the basis coefficients c.

After the collocation method has rendered a solution, the residual function should be computed to evaluate the quality of the approximation. This function measures the difference between the left and right sides of the Bellman equation at arbitrary states s when thee value function is replaced with its approximant (and the optimal basis coefficients c). It would be zero for all states in an exact solution, and will be zero at the collocation nodes for any solution. If the approximation is adequate, the residual function will not depart too far from zero for any arbitrary value of the state in the interval S. If large residuals are obtained, the problem should be re–solved using a different basis–node scheme.

We now consider numerical solutions via collocation methods for several of the models discussed in the last section. In the asset replacement problem of 8.3.1, the stochastic element of the problem is taken to be the replacement cost k:

$$k_{t+1} = \bar{k} + \gamma(k_t - \bar{k}) + \epsilon_{t+1}$$

where ϵ is an $i.i.d.$ normal shock with mean zero and variance σ^2. In the implementation of the collocation method, the shock is discretized using a five–node Gaussian quadrature scheme. Solution to the problem (`demdp01`) demonstrates that for a given asset age, the value of the firm is a downward–sloping function of the replacement cost. The function is kinked at the critical replacement cost, below which the asset is to be replaced. The younger the asset, the greater the value of the firm.

In the economic growth example of 8.4.1, the model is operationalized by assuming a social benefit function $u(c) = c^{1-\alpha}/(1-\alpha)$, with α=0.2, and an aggregate production function $h(x) = x^{\beta}$, with β=0.5. The shock process, which modifies the value of production, is taken to be lognormal with variance σ^2=0.01. The model is coded to incorporate the constraints on the action variable: in this case, to specify that investment x must be non–negative and no greater than s. The model function must also specify the reward function value, the state transition function, and the analytical first and second derivatives of those functions. The lognormal production shock is discretized using a three–node Gaussian quadrature scheme, and a polynomial basis is used on the interval [5,10] for the space of expected wealth. As we see from the graphs (`demdp07`) of a Monte Carlo simulation of this model, the steady state distribution is centered on the value of the certainty–equivalent path of expected wealth, which converges asymptotically to about 7.5 units after 10 years or so.

In the continuous–state mine management problem of 8.4.3, the model is solved using an inverse demand function $p(x) = a_1 - a_2 x$ (that is, linear demand) and a cost of extraction function $c(s, x) = b_1 x - 0.5 b_2 x(2s - x)$ which causes the cost of extraction to rise with the depletion of the mine. Constraints, as in the problem above, are placed on the action space to indicate that the extraction (or "harvest") must be non–negative and no greater than s. Solution of the model indicates a shadow price function with a kink at two units of remaining stock: that is, beyond that point, extraction will never be the optimal strategy. This model is illustrated in `demdp09`.

Finally, in the production–inventory example of 8.4.7, the model is operationalized with quadratic production and cost–of–storage functions: $c(q) = c_1 q + 0.5 c_2 q^2$ and $k(x) = k_1 x + 0.5 k_2 x^2$, respectively. The evolution of the market price is governed by $p_{t+1} = \bar{p} + \rho(p_t - \bar{p}) + \epsilon$ where the latter is an $i.i.d.$ normal shock. The model is considerably more complex in its solution since it is characterized by two states and two actions, requiring two–dimensional grids in each of these spaces. Since the price process is mean–reverting, the optimal inventory policy will be to store nothing if the price is sufficiently high, since it is likely to fall, and the cost of storage will exceed expected appreciation of the good. For sufficiently low prices, it will be economical to hold inventories, since in that instance the expected appreciation of the good will exceed the cost of storage. The value of the firm is an increasing function of both the market price and beginning inventories. With low prices, a simulation (`demdp13`) reveals that the firm will obtain substantial stocks at the outset, but is expected to gradually reduce those stocks over time, reaching a small steady–state mean value of inventories.