Smooth Approximation of the Decision Rule (Cont.)

- An alternative approach to the one we covered the last time is to approximate the decision rule, α, directly, completely bypassing the need to approximate V. Smith (1991).
- Let $\alpha_\theta(s)$ denote a parametric approximation to α. Meaning a smooth mapping that depends on k parameters chosen to best approximate the function. For example α_θ could be given by a polynomial series approximation.
- For example under some conditions we can write

  $$\alpha_\theta(s) = exp[\sum_{i=1}^{k} \theta_i \rho_i(s)], \tag{1}$$

  where the $\rho_i(s)$ are some polynomial basis.
- If there are flexible parameterizations that provide good approximations to a wide class of potential decision rules using a relatively small number of parameters k, then we can approximate the optimal decision rule α by the parametric decision rule $\alpha_{\hat{\theta}}$ where $\hat{\theta}$ is the solution to the k-dimensional static optimization problem:

  $$\hat{\theta} = arg \max_{\theta \in R^k} V_{\alpha_\theta}(s), \tag{2}$$

  and V_{α_θ} is the value function corresponding to the decision rule $\alpha_\theta = (\alpha_{0,\theta}, \cdots, \alpha_{T,\theta})$

  $$V_{\alpha_\theta}(s) \equiv E_{\alpha_\theta}\left[\sum_{t=0}^{T} \beta^t U(s_t, \alpha_{t,\theta}(s_t))|s_0 = s\right]. \tag{3}$$
- For certain Markov Decision Problems the optimal decision rule might be a relatively simple function of s, so the parsimonious characterization α_θ with small k might do a good job in approximating the true decision rule.
- In this cases Smith refers to the approximation as simple rule of thumb.
- The difficulty with this method is that we generally do not know the form of the value function $V_{\alpha_\theta}(s)$ corresponding to the decision rule α_θ. The calculation of this value function is a major source of numerical burden in solving the dynamic programming problem.

- Monte Carlo integration can be used to compute approximations to this value function. I refer you to Rust (1996) and Smith (1991) for extensions of this technique to take into account possible problems with this parsimonious simulator.

Regression-based Interpolation and Monte Carlo Simulation Methods

- Many times we have talked about the curse of dimensionality associated with the solution to the dynamic models we are studying. This can be specially clear when we have to solve multidimensional numerical integration subproblems to evaluate a Bellman operator at a set of points $s \in S$.
- Since randomized Monte Carlo integration success in breaking the curse of dimensionality of the integration problem, these simulation methods might be more effective than deterministic multivariate quadrature methods for solving the integration sub-problem in high-dimensional Markov Decision Problems.
- Keane and Wolpin (1994) use this approach in their algorithm for approximating solutions to finite horizon Discrete Decision Problems (MDP with finite set of choices $a \in A(s)$).
- The Keane-Wolpin algorithm computes the alternative-specific value function $V_t(s,a)$, $t = 0 \cdots, T$, defined by the recursion

  $$V_t(s,a) = u(s,a) + \beta \int \max_{a' \in A(s')} [V_{t+1}(s',a')] p(ds'|s,a), \tag{4}$$

  with the terminal condition $V_T(s,a) = u(s,a)$. The $V_t(s,a)$ functions are related to the usual value functions $V_t(s)$ by the identity:

  $$V_t(s) = \max_{a \in A(s)} [V_t(s,a)]. \tag{5}$$

  To implement the algorithm one needs to specify a grid over the state space S, for example $(s_1, \cdots, s_N)$. At each point s_j on the grid and for each action $a \in A(s_j)$ one draws N realizations from the transition density $p(ds'|s_j,a)$. Denote these realizations $(\tilde{s}_{1ja}, \cdots, \tilde{s}_{Nja})$ to emphasize the fact that separate realizations are drawn for each conditioning pair (s_j, a).

- Then the Keane-Wolpin estimate of $V_t(s_j,a)$ is given by

$$\hat{V}_t(s_j,a) = u(s_j,a) + \frac{\beta}{N}\sum_{i=1}^{N} \max_{a' \in A(\tilde{s}_{ija})} \left[\hat{V}_{t+1}(\tilde{s}_{ija'},a')\right], j = 1,\cdots,N. \qquad (6)$$

  Notice that $\hat{V}_t$ is only defined at the pre-assigned grid points $(s_1,\cdots,s_N)$ an not at other points $s \in S$. Since the realizations $(\tilde{s}_{1ja},\cdots,\tilde{s}_{Nja})$ will generally not fall on the pre-assigned grid points, they fit a linear regression using the N simulated values $y \equiv (\tilde{V}_t(s_1),\cdots,\tilde{V}_t(s_N))$ as an N by one dependent variable, and various functions of the components of the estimated expected value function $EV_{t+1}(s,a) \equiv \int V_{t+1}(s')p(ds'|s,a)$ evaluated at the grid points $(s_1,\cdots,s_N)$ to form the N *by* N matrix of regressors, X.

- The OLS coefficient estimates $\hat{\beta} = (X'X)^{-1}X'y$ are then used to generate estimates of $\tilde{V}_t(s)$ at arbitrary points $s \in S$, outside of the pre-assigned grid.

- In principle we could use any non-parametric regression technique, or any interpolation methods to do this step.

- Notice that in equation (6) it is assumed that some sort of interpolation procedure has already been carried out on the stage $t+1$ value function $\hat{V}_{t+1}$ so that it can be evaluated at all the random points $\tilde{s}_{ija}$, $i = 1,\cdots,N$, that were drawn at the stage t grid point (s_j,a).

- They conclude that their approximation method ameliorates the curse of dimensionality problem, obtaining approximate solutions for problems with otherwise intractably large state space.

- However, their technique cannot avoid the curse of dimensionality of the approximation subproblem, as was shown by Novak (1988) and others.

- In a series of papers by Keane and Wolpin and Eckstein and Wolpin they applied this solution technique to the estimation of models of high school attendance, and work during high school