Economics 250c
Fall 2008

We continue a presentation of dynamic discrete choice problems with extreme value errors. We will cover

1. Stange's model of college progression (another simple DDC model)
   Kevin Stange. "An Empirical Examination of the Option Value of College Enrollent" Unpublished paper.
2. The Hotz-Miller approach

Prologue - Learning (Bayesian updating)

In many applications a decision maker is uncertain about the true value of some key parameter, and receives new information over time about the value of the parameter. In Stange's case the key parameter is A, which reflects the "difficulty" of college-level work. The natural way to model this class of problems is using Bayesian updating with conjugate priors. The classic reference is de Groot, 1970.

Normal learning.

True state variable is η, with $-\infty < \eta < \infty$. Prior on η is $N(m_0, 1/H_0)$. The observed signal is $s = \eta + \epsilon$, with $\epsilon \sim N(0, 1/h)$, independent of η. It can be shown that posterior for η is

$$N\left(\frac{H_0 m_0 + hs}{H_0 + h}, \frac{1}{H_0 + h}\right).$$

With a sequence of observations $s_t = \eta + \epsilon_t$, with $\epsilon_t \sim N(0, 1/h)$, the posterior after the 1st observation has mean m_1 and precision H_1 given by these formulas. Preceding sequentially, the posterior after the t^{th} observation, conditional on the mean and precision after the $(t-1)^{st}$, is normal with mean and precision

$$\begin{aligned} m_t &= \frac{H_{t-1} m_{t-1} + h s_t}{H_{t-1} + h} = \frac{H_0 m_0 + h \sum_{k=1}^t s_k}{H_0 + th} \\ H_t &= H_{t-1} + h = H_0 + th \end{aligned}$$

Note that

$$\begin{aligned} m_t &\rightarrow \frac{1}{t}\sum_{k=1}^{t} s_k \quad \text{the mean of the signals up to period } t \\ H_t &\rightarrow th \quad \text{so } \frac{1}{H_t} \rightarrow \frac{1}{t} var[\epsilon_t] \quad \text{the variance of the mean of the signals up to } t \end{aligned}$$

This formula has had many applications in labor economics, e.g. models of learning about match quality.

Beta-Bernoulli

Suppose y_t is distributed as a Bernoulli with $P(y_t = 1) = p$. The conjugate prior for p is $Beta(\alpha, \beta)$. For $p \sim Beta(\alpha, \beta)$:

$$\begin{aligned} E[p] &= \frac{\alpha}{\alpha + \beta}, \\ var[p] &= \frac{\alpha\beta}{(\alpha + \beta)^2 (1 + \alpha + \beta)} \end{aligned}$$

The density is $f(p) = \frac{\Gamma(\alpha)}{\Gamma(\alpha)\Gamma(\beta)} p^{\alpha-1}(1-p)^{\beta-1}$. Note that $Beta(1,1) = U(0,1)$. The posterior for p, given a draw y_1 is $Beta(y_1+\alpha, 1-y_1+\beta)$. Applying this sequentially, the posterior after t realizations with S_t successes is $Beta(S_t+\alpha, t-S_t+\beta)$, implying the posterior mean and variance are

$$
\begin{aligned}
E[p|S_t] &= \frac{S_t+\alpha}{t+\alpha+\beta} \to \frac{1}{t}S_t \\
var[p] &= \frac{(S_t+\alpha)(t-S_t+\beta)}{(t+\alpha+\beta)^2(1+\alpha+\beta+t)} \to \frac{\frac{1}{t}S_t \times (1-\frac{1}{t}S_t)}{t}
\end{aligned}
$$

A nice application of this class of learning models is to problems of the form "waiting for a prize that will arrive with probability p." If the model is formulated so the agent "opts out" of the wait when p is low, then optimal behavior is to wait until n unsuccessful draws, then opt out. In his thesis, L. Katz applied this idea to the behavior of workers on temporary layoff, who have to decide whether to continue waiting for recall, or start looking for a new job. McCall applied it to a problem of hierarchical search over occupations and then specific jobs within an occupation. In each case the formula delivered an expression for the optimal time to wait before "bailing out".

1. Stange's Model

Stange's model is motivated by the observation that the Mincerian "returns" to the first 3 years of college are very low, while the return to a 4-year degree are very high. Why do so many people attend college for 1-3 years and then drop out? His explanation is that they have to go to college to learn if they have the ability to complete college. The model explains the demand for junior college as a cheap option for learning about ability.

The model has two very important simplifying assumptions. First, it is assumed that school-leaving is irreversible. (More general models, e.g. Keane and Wolpin JPE 1997 add a "fixed cost" of returing to school once you drop out – they estimate this is very large, ~\$10,000). Second, it is assumed that people know their expected discounted income at each of the 5 "exit nodes" (after highschool, after 1 year of college, ... after 4 years of college). Stange imputes this number, called $Income_1, Income_2, ...Income_5$, using data for an earlier cohort in the NLSY. The estimates are conditional on exit node, and some characteristics: race dummies, region dummies, and a quadratic in {HSGPA, AFQT, Parental Education}.

In the absence of learning, Stange's model is very simple. At each of the 4 decision points ($t = 1$ for the end of high school, $t = 2$ for the end of 1st year of college, ... $t = 4$ for the end of third year of college) the student decides whether to continue or stop. At decision point t the value of entering the labor market is

$$V_t^w = Income_t + \epsilon_t^w,$$

where ϵ_t^w/τ is EV1. The flow value of continuing in school for the next year is

$$\overline{u}_t^s + \ \epsilon_t^s,$$

where ϵ_t^s/τ is EV1. Adding the discounted expected value of entering the next node, we have

$$V_t^s = \overline{u}_t^s + \ \epsilon_t^s + \beta E_t[V_{t+1}^s].$$

At decision point t the optimal choice involves a comparison between V_t^w and V_t^s.

Learning

The innovation in Stange's model is the introduction of another variable, A, that represents ability and affects the cost of going through college. Specifically, Stange assumes that

$$\overline{u}_t^s = \alpha_m + \alpha_A E_t[A] - \alpha_d Dist_t - Tuition_t + \alpha_{2year}\ 1(\text{enrolled in 2-year})$$

where α_m is a random effect. This specification is somewhat ad hoc: one could imagine a model where the relation between effort, latent ability, and grade outcomes is modelled explicitly. Students learn A through their college grades, with a prior as of the end of high school that depends on HSGPA, AFQT, and parental education:

$$E_1[A] = \gamma_m + \gamma_g HSGPA + \gamma_t AFQT + \gamma_p ParentEd$$

Again, γ_m is a random effect. The pair (α_m, γ_m) are assumed to have a discrete bivariate distribution with M (=3) points of support. He assumes that

$$\begin{aligned} E_t[A] &= k_t\gamma + (1-k_t)GPA_{t-1}, \quad \text{where} \\ GPA_{t-1} &= \frac{1}{t-1}\sum_{k=1}^{t-1} g_k \ \text{ is the cumulative GPA up to t-1,} \end{aligned}$$

and g_k = GPA in year k of college. In principle, if $g_k = A + e_k$, and both A and the "grade shocks" e_k are normally distributed, then k_t (which Stange calls γ_{Xt}) should look like the term in a normal-learning model. He does not impose that, but treats $\{k_t\}$ as parameters to be estimated.

With this setup it is straightforward to solve backward from the final decision node. At node 4 (end of 3rd year), the student has observed g_1, g_2, g_3, and has to compare:

$$\begin{aligned} \text{value of dropout} &= V_4^w = Income_4 + \epsilon_4^w \\ \text{value of continuing} &= V_4^s = \overline{u}_4^s + \epsilon_4^s + Income_5 + E[\epsilon_5^w] \\ &= \alpha_m + \alpha_A E_4[A|GPA_3, X] - \alpha_d Dist_4 - Tuition_4 + \epsilon_4^s + Income_5 + c. \end{aligned}$$

where c=$E[\epsilon_5^w]$. The $E\max$ of this, conditional on GPA_3, X is

$$E\max[V_4^w, V_4^s|GPA_3, X]$$

which has the usual form, assuming ϵ_4^w and ϵ_4^w are scaled EV1's. We also can form a standard logit-type probability for dropping out at node 4, conditional on the information up to that point:

$$P(dropout\ at\ node\ 4|g_1, g_2, g_3, X) = \frac{\exp(Income_4)}{\exp(Income_4) + \exp(\overline{u}_4^s + Income_5 + c)}.$$

At node 3 (end of 2nd year), the student has observed GPA_2 and has to compare:

$$\begin{aligned} \text{value of dropout} &= V_3^w = Income_3 + \epsilon_3^w \\ \text{value of continuing} &= V_3^s = \overline{u}_3^s + \epsilon_3^s + E_3 E\max[V_4^w, V_4^s|GPA_3, X] \\ &= \alpha_m + \alpha_A E_3[A|GPA_2, X] - \alpha_d Dist_3 - Tuition_3 \\ &\quad + E_3 E\max[V_4^w, V_4^s|GPA_3, X] + \epsilon_3^s \end{aligned}$$

Note that the calculation of $E_3 E\max[V_4^w, V_4^s|GPA_3, X]$ requires a probability distribution for GPA_3 conditional on GPA_2 (since g_3 is not known at node 3). Stange uses the usual "Rust" assumption that

$$E_3 E\max[V_4^w, V_4^s|GPA_3, X] = \int E\max[V_4^w, V_4^s|GPA_3, X]\pi(GPA_3, X|GPA_2, X)$$

where π gives the "transition probabilities" from GPA_2 to GPA_3. To simplify, Stange assumes that cumulative GPA lies on a discrete grid. Assuming g_3 is normally distributed allows him to use simple expressions for the probabilites of alternative values of GPA_3 conditional on GPA_2. With this calculation in hand it is then possible to calculate P(dropout at node $3|GPA_2, X$). The calculations for nodes 2 and 1 are similar, and require calculations of the expressions

$$\begin{array}{l} E_2 E\max[V_3^w, V_3^s|GPA_2, X] \text{ and} \\ E_1 E\max[V_2^w, V_2^s|GPA_1, X]. \end{array}$$

The model delivers a likelihood for the joint distribution of GPA outcomes and dropout decisions. The calculations are pretty simple because the model gives the conditional probabilities of dropping out at each node, and the conditional probabilites for GPA_t conditional on GPA_{t-1}. With mass-point mixing the likelihood is evaluated at each pair of mass points, then averaged across the mass points with weights that are jointly estimated.

2. The Hotz-Miller Approach (introduction)

In lecture 6 we considered a choice problem with J choices, where $u_j = v_j + \epsilon_j$, with v_j a set of functions whose form is known and $(\epsilon_1, \epsilon_2, ...\epsilon_J) \sim F(\epsilon_1, \epsilon_2, ...\epsilon_J)$. Choice 1 is selected when $v_1 + \epsilon_1 > v_k + \epsilon_k$, or $\epsilon_k < v_1 - v_k + \epsilon_1$ (for all $k = 2, ..J$), which has probability

$$\begin{aligned} p_1 &= \int_{-\infty}^{\infty}\int_{-\infty}^{v_1-v_2+\epsilon_1} ... \int_{-\infty}^{v_1-v_J+\epsilon_1} f(\epsilon_1, \epsilon_2, ...\epsilon_J)d\epsilon_2...d\epsilon_J\ d\epsilon_1, \\ &= \phi_1(v_1 - v_2, v_1 - v_3, ...v_1 - v_J). \end{aligned}$$

Similarly for choices $2, 3, ...J$, we can write

$$p_j = \phi_j(v_j - v_1, v_j - v_3, ...v_j - v_J).$$

The HM "representation theorem" is that we can invert this mapping to get

$$v_j - v_k = \psi_{jk}(p),\ where\ \ p = (p_1, ...p_J).$$

For example, in the simple MNL

$$p_j = \frac{e^{v_j}}{\sum_k e^{v_k}},$$

so (as noted by Berry, 1994):

$$\log p_j - \log p_k = v_j - v_k.$$

Thus, in the basic MNL case, $\psi_{jk}(p) = \log(p_j/p_k)$. Given the mapping from the $p's$ to the $v's$, we noted that the "selection bias" terms can as be written in terms of the choice probabilities:

$$\begin{aligned} E[\epsilon_1|d_1 = 1] &= \frac{\int_{-\infty}^{\infty}\int_{-\infty}^{v_1-v_2+\epsilon_1} ... \int_{-\infty}^{v_1-v_J+\epsilon_1} \epsilon_1 f(\epsilon_1, \epsilon_2, ...\epsilon_J)d\epsilon_2...d\epsilon_J d\epsilon_1}{P(d_1 = 1|v_1..v_J)} \\ &= w_1(\psi(p)) \end{aligned}$$

for some function w_1. Likewise for the other choices:

$$E[\epsilon_j | d_j = 1] = w_j(\psi(p)).$$

For example, recall from lecture 6 that for the basic MNL:

$$E[\epsilon_j | d_j = 1] = \gamma - \log p_j$$

where $\gamma =$ Euler's constant.

Now consider a finite horizon discrete dynamic programming problem with a state variable s_t, where in each period the agent chooses from a set $\{d_1, ...d_J\}$ and gets within-period (flow) utility $u_j(s_t) + \epsilon_{jt}$ from choice j if the state is s_t. Let $d_t = (d_{1t}...d_{Jt})$ denote the vector of indicators for the choices in t (one of these is 1 and the rest are 0). Assume that the vector of shocks in period t, ϵ_t, is i.i.d. distributed over time, and let $F_k(s_{t+1}|s_t)$ represent the d.f. for s_{t+1} conditional on s_t and $d_{kt} = 1$. The individual's objective is:

$$\max \quad E\left(\sum_{t=1}^{T} \sum_{j=1}^{J} \beta^t d_{jt} [u_j(s_t) + \epsilon_{jt}] \ \mid \ s_t \right),$$

where expectations are taken with respect to the joint distribution of future states and the ϵ_t's.

Let $d^0_{jt}(s_t, \epsilon_t)$ represent the optimal decision rule in period t conditional on s_t and ϵ_t. Define $V(s_t)$ as the expected payoff associated with being in state s_t, assuming optimal choices from period t onward:

$$V(s_t) = E\left(\sum_{\tau=t}^{T} \sum_{j=1}^{J} \beta^{\tau - t} d^0_{j\tau}(s_\tau, \epsilon_\tau)[u_j(s_\tau) + \epsilon_{j\tau}] \ \mid s_t \right).$$

Finally, define

$$(*) \qquad \begin{aligned} v_k(s_t) &\equiv u_k(s_t) + \beta E[V(s_{t+1})|s_t] \\ &= u_k(s_t) + E\left(\sum_{\tau=t+1}^{T} \sum_{j=1}^{J} \beta^{\tau-t} d^0_{j\tau}(s_\tau, \epsilon_\tau)[u_j(s_\tau) + \epsilon_{j\tau}] \ \mid s_t \right) \\ &= u_k(s_t) + \sum_{\tau=t+1}^{T} \sum_{j=1}^{J} \beta^{\tau-t} p_j(s_\tau)[u_j(s_\tau) + E[\epsilon_{j\tau} | \ d_{j\tau} = 1, s_\tau]] \end{aligned}$$

where $p_j(s_\tau)$ represents the "conditional choice probability" for choice j in period τ when the state is s_τ. Note that $v_k(s_t) + \epsilon_{kt}$ is the value of choosing alternative k in period t when the state is s_t.

We can apply the HM representation theorem to express the differences in the $v_k(s_t)$ functions (for a given state) in terms of the choice probabilies:

$$v_j(s_t) - v_k(s_t) = \psi_{jk}(p(s_t)),$$

where $p(s_t) = (p_1(s_t), ...p_J(s_t))$. We can also write the selection bias terms as functions of the choice probabilities:

$$E[\epsilon_{kt} | d_{kt}, s_t] = w_k(\psi(p(s_t)).$$

Now suppose we knew the form of the selection corrections. Then we could rewrite the last line of equation $(*)$ as:

$$(**) \qquad v_k(s) = u_k(s) + E\left(\sum_{\tau=t+1}^{T} \sum_j \beta^{\tau-t} p_j(s_\tau)[u_j(s_\tau) + w_j(\psi(p(s_\tau))] \mid s_t = s, d_{kt} = 1\right).$$

This expresses everything in terms of the flow payoff functions $u_k(s)$, and the conditional choice probabilities $p_j(s_\tau)$. The idea of the HM approach is to start with estimates of the conditional choice probabilities, and choose distributions for the error terms so that $w_j(\psi(p(s_\tau))$ can be computed. Then using $(**)$ it is possible to get expressions for the conditional valuation of choice k in state s. Intuitively, what we are doing is "filling in" the expression $\beta E[V(s_{t+1})|s_t]$ with an average of the payoffs that the individual will receive in each future state, weighting the payoff associated with each choice by the estimated probability that this is the optimal choice, and adjusting for the selection bias term.