Economics 250c
Fall 2008

In this lecture we discuss estimation of discrete dynamic choice models using conditional choice probabilities, as first proposed by Hotz and Miller (1993). We then consider in some detail the Rust 1987 paper, and the application of the HM approach to this problem.

A very good reference:

Victor Aguirregabiria and Pedro Mira. "Dynamic Discrete Choice Structural Models: A Survey." Available at SSRN.

Recap from the end of last lecture:

For a random utility model with J choices, and payoffs $u_j = v_j + \epsilon_j$, HM showed there is a inversion mapping

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

For example, in the simple MNL

$$\begin{aligned} p_j &= \frac{e^{v_j}}{\sum_k e^{v_k}} \\ &\Rightarrow \ \log p_j - \log p_k = v_j - v_k \end{aligned}$$

Thus, for the MNL, $\psi_{jk}(p) = \log(p_j/p_k)$. It is also possible to write the "section bias" $E[\epsilon_j | d_j = 1]$ in terms of the vector of differences $v_j - v_k$. Thus, in general we have

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

E.g., for the MNL:

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

where $\gamma =$ Euler's constant.

Now consider a finite state discrete dynamic choice 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 $u_j(s_t|\theta) + \epsilon_{jt}$ from choice j if the state is s_t, where θ are a set of unknown parameters. 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 that state transitions are given by functions $P(s_{t+1}|s_t, d_t)$.

Note that this setup satisfies the 2 critical assumptions of the "Rust approach"

additivity between $u_j(s_t)$ and ϵ_{jt} (with ϵ_{jt} distributed over $(-\infty, +\infty)$)

"conditional independence" $p(\epsilon_t, s_t | s_{t-1}, \epsilon_{t-1}, d_{t-1},) = p(\epsilon_t|s_t) p(s_t|d_{t-1}, s_{t-1})$

We can write the individual's objective starting at period 1 as

$$\max \ \ E\left(\sum_{t=1}^{T} \sum_{j=1}^{J} \beta^t d_{jt}[u_j(s_t|\theta) + \epsilon_{jt}] \ \ | \ \ s_1 \right),$$

where expectations are taken with respect to the joint distribution of future states and the ϵ_t's, assuming that the individual behaves optimally in the future. (To actually write out the expectation is quite an effort in notation). Let $d^0_{jt}(s_t, \epsilon_t)$ represent the optimal decision rule in period t conditional on s_t and ϵ_t. Define $\overline{V}(s_t)$ (the integrated value function) as the expected payoff associated with being in state s_t, assuming optimal choices from period t onward:

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

Finally, define the choice specific value functions

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

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.

Write the selection bias terms as functions of the choice probabilities:

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

This allows us to rewrite the last line of equation $(*)$ as:

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

If we can estimate the functions $p(s_\tau)$, and we know the $w_j(\psi(p(s_\tau))$ functions (as is true for MNL and GEV models) then we express $v_k(s_t)$ in terms of current and future flow utility functions and the choice probabilities. Call the estimated choice probabilites $\widehat{p}(s_\tau)$, and let

$$\widehat{v}_k(s_t|\theta) = u_k(s_t|\theta)+E\left(\sum_{\tau=t+1}^{T}\sum_{j=1}^{J}\beta^{\tau-t}\widehat{p}_j(s_\tau)[u_j(s_\tau|\theta)+w_j(\psi(\widehat{p}(s_\tau))] \mid s_t=s, d_{kt}=1\right).$$

(Note that v_k and $\overline{V}$ both depend on θ but until this last equation the dependence was implicit). Given a value for θ the implied probability that choice k is selected when the state is s_t is

$$p_k(s_t) = \frac{\exp(\widehat{v}_k(s_t|\theta))}{\sum_j \exp(\widehat{v}_j(s_t|\theta))}$$

HM proposed to estimate θ by applying GMM to the deviations of the estimated probabilites from the predicted values (which themselves depend on the probabilities).

Aguirregabiria and Mira (Econometrica, 2002) suggested the alternative idea of estimating θ by maximizing the likelihood of the observed sequence of choices "as if" $v_k(s_t)$ were the true payoff function for choice k in state, i.e., by maximizing the "pseudo likelihood"

$$Q = \sum_{i=1}^{n} \sum_{t=1}^{T} \log \left(\frac{\exp(\widehat{v}_{k(i,t)}(s_t)}{\sum_k \exp(\widehat{v}_k(s_t))} \right)$$

where $k(i,t)$ is the actual choice made by individual i in period t.

Provided that the estimated conditional choice probabilities are $\sqrt{n}-$consistent, A-M showed the estimates of θ are asymptotically equivalent to estimates obtained by a 2-step "nested fixed point" algorithm. (see below for what this is).

A-M's pseudo-likelihood method is especially easy in the special case where $u_j(s_t|\theta) = z(s_t, j)\theta$. In this case, looking at $(**)$ you see that θ can be taken outside the sum of the terms involving $z(s_\tau, j)\theta$. Thus, $\widehat{v}_k(s_t|\theta)$ can be written as

$$\widehat{v}_k(s_t|\theta) = \widetilde{z}(s_t, k)\theta + \widetilde{w}_k(s_t)$$

where

$$\widetilde{z}(s_t, j) = E\left(\sum_{\tau=t+1}^{T} \sum_{j=1}^{J} \beta^{\tau-t} \widehat{p}_j(s_\tau) z(s_\tau, j) \mid s_t = s, d_{kt} = 1 \right)$$

and

$$\widetilde{w}_k(s_t) = E\left(\sum_{\tau=t+1}^{T} \sum_{j=1}^{J} \beta^{\tau-t} \widehat{p}_j(s_\tau) w_j(\psi(\widehat{p}(s_\tau)) \mid s_t = s, d_{kt} = 1 \right)$$

(the expectations are over the transition probabilities from $s_t \to s_{t+1} \to s_{t+2}...$). Note that the sums $\widetilde{z}(s_t, j)$ and $\widetilde{w}_k(s_t)$ only have to be calculated once! The function $\widetilde{z}(s_t, j)$ is a discounted, probability-weighted average of the future $z(s_\tau, j)'s$ that the decision maker will get to if he/she makes the j^{th} choice in period t with state s_t.

A-M 2002 also propose an iterative procedure:

given an estimated set of choice probabilities, maximize the pseudo likelihood to estimate θ

given θ, solve the dynamic problem and get a new set of choice probabilites to construct $\widetilde{z}$ and $\widetilde{w}_k$.

maximize the pseudo likelihood to get a new estimate of θ

They show that this converges to the solution from the NFP algorithm applied to the same problem. The iterative solution does not yield any gain in asymptotic efficiency but has better small sample properties.

Rust (1987)

Let's change notation slightly to more closely match Rust, and consider a discrete time, infinite horizon problem with a finite set of states denoted by x_t and a fixed set of choices denoted by $d_1, ...d_J$. The flow payoff in period t when the agent makes choice d_j is

$$u(x_t, d_j|\theta_1) + \epsilon_{jt}$$

and the transition probabilites are

$$p(x_{t+1}|x_t, d_t, \theta_3)$$

where θ_1 and θ_3 are parameters (don't ask what happened to θ_2). Temporarily dropping the $\theta's$ denote the value function (in the notation of lecture 7) as

$$\begin{aligned} V(x_t, \epsilon_t) &= \max_d \quad u(x_t, d) + \epsilon(d) + \beta EV(x_{t+1}, \epsilon_{t+1}) \\ &= \max_d \quad u(x_t, d) + \epsilon(d) + \beta \sum_{x_{t+1}} \left[\int_{\epsilon_{t+1}} V(x_{t+1}, \epsilon_{t+1}) dF(\epsilon_{t+1}) \right] p(x_{t+1}|x_t, d) \end{aligned}$$

The expected value of V (the integrated value function) is

$$\overline{V}(x_t) = \int_{\epsilon_t} \left[\max_d \quad u(x_t, d) + \epsilon(d) + \beta \sum_{x_{t+1}} \overline{V}(x_{t+1}) p(x_{t+1}|x_t, d) \right] dF(\epsilon_t).$$

Rust's crucial insight is that $\overline{V}(x_t)$ looks just like the $E\max$ for a problem with payoffs $v(x_t, d) + \epsilon(d)$ associated with choices d, where

$$v(x_t, d) = u(x_t, d) + \beta \sum_{x_{t+1}} \overline{V}(x_{t+1}) p(x_{t+1}|x_t, d)$$

is the "choice specific value function". Moreover, the probability that choice d is made (assuming the agents observe $\epsilon(d)$ when they get to period t, but we as econometricians do not) is

$$P(d_{jt} = 1|x_t) = \frac{\exp(v(x_t, j))}{\sum_k \exp(v(x_t, k))}.$$

So the way to solve the infinite horizon choice problem is to find a value function such that

$$\overline{V}(x_t) = E \max_d \{u(x_t, d) + \epsilon(d) + \beta \sum_{x_{t+1}} \overline{V}(x_{t+1}) p(x_{t+1}|x_t, d)\}$$

In the case where the $\epsilon_t's$ are EV1,

$$\overline{V}(x_t) = \gamma + \log \left(\sum_d \left[\exp(u(x_t, d) + \beta \sum_{x_{t+1}} \overline{V}(x_{t+1}) p(x_{t+1}|x_t, d) \right] \right).$$

This functional equation is a contraction mapping, so, given $u(x_t, d)$ and $p(x_{t+1}|x_t, d)$ it is possible to solve for$\overline{V}(x_t)$ by naive iteration. Start with a "trial function" $\overline{V}^1(x_t)$ (which is really just a list of numbers for each possible value of the state variable). Then iterate:

$$\overline{V}^{i+1}(x_t) = \gamma + \log \left(\sum_d \left[\exp(u(x_t, d) + \beta \sum_{x_{t+1}} \overline{V}^i(x_{t+1}) p(x_{t+1}|x_t, d) \right] \right)$$

and stop when successive iterations are very close at all values of the state space.

Rust's 2-step NFP algorithm:
step 1: estimate θ_3 by fitting models for the observed transition rates $p(x_{t+1}|x_t, d_t, \theta_3)$.
step 2: estimate θ_1 in an iterative inner/outer loop procedure
in the inner "contraction mapping" loop, solve for $\overline{V}(x_t)$ and $P(d_{jt} = 1|x_t)$.
use this to construct a partial log likelihood $\sum_t \log P(d_{jt} = 1|x_t)$

in the outer loop, choose over θ_1's until convergence

(as discussed in A-M's review paper, there is a sophisticated way to do this that gets the derivatives of the log likelihood w.r.t. θ_1 "for free" without having to do numerical derivatives).

Some details of Rust's example:

Rust applies this idea to data on engine replacements for a fleet of 162 buses owned by the Madison Metropolitan Bus Company over the period from December 1974 to May 1985. The data consist of monthly observations on each buses oddometer reading, plus an indicator for whether the engine was subject to a major overhall/replacement.

The state space has 90 elements, $x_t \in \{0, 1, ...89\}$, with $x_t = j$ implying that the oddometer reading is in the j^{th} bin (of width 5000 mi). The choice set is $d \in \{0, 1\}$, where $d = 1$ is a "renewal" (replace the engine) and $d = 0$ is don't replace and perform normal maintenance. The payoffs are

$$\begin{aligned} u(x_t, 0) &= -c(x_t, \theta_1) \\ u(x_t, 1) &= -c(0, \theta_1) - RC \end{aligned}$$

where $c(x_t, \theta_1) = \theta_{11} x_t$ in the simplest model. (In the more complicated models c is quadratic or cubic). The transition probabilites are

$$\begin{aligned} p(x_{t+1}|x_t, 0) &= \theta_{31} \;\; if \quad x_{t+1} = x_t \\ &= \theta_{32} \;\; if \quad x_{t+1} = x_t + 1 \\ &= \theta_{33} \;\; if \quad x_{t+1} = x_t + 2 \\ &= 0 \quad otherwise. \end{aligned}$$

In other words, the oddometer can stay the same, rise by 1 bin, or by 2 bins. With this setup the contraction mapping iteration is

$$\begin{aligned} \overline{V}^{i+1}(x_t) &= \gamma + \log\{\exp\left[(-c(x_t, \theta_1) + \beta\theta_{31}\overline{V}^i(x_t) + \beta\theta_{32}\overline{V}^i(x_{t+1}) + \beta\theta_{33}\overline{V}^i(x_{t+2})\right] \\ &\quad + \exp\left[(-c(0, \theta_1) - RC + \beta\theta_{31}\overline{V}^i(0) + \beta\theta_{32}\overline{V}^i(1) + \beta\theta_{33}\overline{V}^i(2)\right]\}. \end{aligned}$$

This converges very fast.

Now lets use the HM approach. A first observation is that if we have

$$v(x_t, d) = u(x_t, d) + \beta \sum_{x_{t+1}} \overline{V}(x_{t+1}) p(x_{t+1}|x_t, d).$$

and d=1 is a "renewal" then $u(x_t, 1) = u(0, 1)$ and $p(x_{t+1}|x_t, 1) = p^*(x_{t+1}|1)$ independent of x_t. (Under Rust's assumptions, after a renewal the state variable can only be at 0, 1 or 2; e.g., $p^*(x_{t+1} = 0|1) = \theta_{31}$). Thus

$$v(x_t, 1) = u(0, 1) + \beta \sum_{x_{t+1}} \overline{V}(x_{t+1}) p^*(x_{t+1}|1) = u(0, 1) + \beta V^*.$$

The value of a renewal does not depend on the state you enter it from (which is what it means to be a "renewal"). A second observation is that

$$\begin{aligned} \overline{V}(x_{t+1}) &= \gamma + \log[\exp(v(x_t, 0)) + \exp(v(x_t, 1))] \\ &= \gamma + \log[\exp(v(x_t, 0)) + \exp(u(0, 1) + \beta V^*)] \end{aligned}$$

Third, note that

$$P(renew \text{ at } \; t+1) = P_1(x_{t+1}) = \frac{\exp(u(0,1)+\beta V^*)}{\exp(v(x_t,0)) + \exp(u(0,1)+\beta V^*)},$$

implying that

$$\log P_1(x_{t+1}) = u(0,1) + \beta V^* - \log[\exp(v(x_t,0)) + \exp(u(0,1)+\beta V^*)].$$

Combining this with the expression for $\overline{V}(x_{t+1})$, we have

$$\overline{V}(x_{t+1}) = \gamma + u(0,1) + \beta V^* - \log P_1(x_{t+1}).$$

{Aside: this shows a general result that the emax can always be written as the value of 1 of the alternatives, plus γ minus the log of the probability of that alternative. In a renewal problem, it is convenient to use this idea to express the emax for any value of the state variable in terms of the renewal probability from that state.}

Plugging in to the equation for $v(x_t,0)$, we get

$$\begin{aligned} v(x_t,0) &= u(x_t,0) + \beta \sum_{x_{t+1}} (\gamma + u(0,1) + \beta V^* - \log P_1(x_{t+1}))\, p(x_{t+1}|x_t,0) \\ &= u(x_t,0) + \beta(\gamma + \beta V^*) + \beta \sum_{x_{t+1}} (u(0,1) - \log P_1(x_{t+1}))\, p(x_{t+1}|x_t,0). \end{aligned}$$

This says that we can calculate the full value of not renewing at state x_t using only information on $u(x_t,0)$, $u(0,1)$, the transition probabilities, and the conditional renewal probabilities. We can use the expressions for $v(x_t,1)$ and $v(x_t,0)$ to get the probability of renewal in each state in terms of the primitives, then fit a pseudo likelihood as recommended by AM.