Dynamic *and* Stochastic Model of Industry

- Agenda motivated by facts about churning in industries
  - Some firms grow in same industry where other firms decline
  - In same industry, new firms enter while incumbents exit.

- Beyond the above fact, agenda produces a framework that is in principle consistent with lots of different kinds of behavior
  - Model with parameters of cost and demand and how they change over time
  - Model includes oligopolistic interactions

- Example of what to do with this this?
  - Examine effects of policies (run counterfactuals), e.g. mergers, environmental policies,...

Start Simple: Single Agent Problem

- John Rust Bus Engire Replacement problem. (More generally think of a firm replacing a machine. For every replacement, there is an exit of the incumbent machine and entry of a new machine.)

- Time $t = 1, 2, ...$

- Actions $a_t \cancel{=} \overset{t}{A}$

  - $a_t = 0$ means keep current machine

  - $a_t = 1$ means replace.

- State variables at time t
  - Condition of incumbent machine ω_t. Let this be an integer, $0 \leq \omega \leq \bar{\omega}$. When new it equals $\bar{\omega}$ ("observed" by us as well as agent)
  - A utility shock to each choice ("unobserved state variable" by us (agent making decision sees this)
    - $\varepsilon_{t,0}$ utility shock to $a_t = 0$
    - $\varepsilon_{t,1}$ utility shock to $a_t = 1$
    - Standard to assume i.i.d. If we further assume that it is Type I extreme value we obtain considerable analytic tractability

- Transition probabilities $\Pr(\omega_{t+1}|a_t, \omega_t)$ for incumbent machine
  - If replace machine then $\omega_{t+1} = \bar{\omega}$
  - It don't replace machine:
    * One possibility is deterministic decay, $\omega_{t+1} = \omega_t - 1$, if $\omega_t > 0$ and $\omega_{t+1} = 0$, if $\omega_t = 0$.
    * In general, could make the transition stochastic (and including even random improvements in condition).

- Payoffs (other than utility shock mentioned above
  - $\pi_{\omega,0}$ if keep current machine
  - $\pi_{\omega,1}$ if replace
  - For example $\pi_{\omega,0} < \pi_{\omega+1,0}$, and $\pi_{\omega,1} = \pi_{\omega,0} - \eta$ (where η is the cost of replacement)

- Let θ be a vector of parameters
  - Includes payoffs, $\pi_{\omega,a}$ and the transition Pr $(\omega'|\omega, a))$
  - Normalize the extreme value distribution to the standardized value. This is w.l.o.g. since can rescale by $\lambda > 0$

$$\begin{aligned} \pi'_{\omega,a} &= \lambda\pi_{\omega,a} \\ \varepsilon'_{it} &= \lambda\varepsilon_{it} \end{aligned}$$

    the λ factors out and all decisions are the same. Suppose directly observe the replacement cost R(e.g. $R =$\$10,000). Then can define $\eta = \alpha R$, and this gives us the utility weight on money, given the normalization. (Equivalently, could normalize $\alpha = 1$, and then include a parameter to rescale the ε_{it}.)
  - Can thinking of this as an entry model, with η as the entry cost. Perhaps we don't observe it. But maybe we

observe π_ω for two of the ω states ω, e.g. $\pi_{\bar{\omega}} > \pi_{\bar{\omega}-1}$. (Or the dollar value associated with this state which we can multiply by α). In this way, we can back out the entry cost. This is the big idea of the approach. Can use revealed preferences to infer switching costs.

Facts about Logit Error Structure

- Extreme Value Type I. ε has CDF

$$\Pr(\varepsilon < c) = F(c) = \exp(-\exp(-c))$$

- Suppose i.i.d. draws ε_i and $\varepsilon_{i'}$, then distribution of the difference is logistic

$$F(\varepsilon_i - \varepsilon_{i'}) = \frac{\exp(-(\varepsilon_i - \varepsilon_{i'}))}{1 + \exp(-(\varepsilon_i - \varepsilon_{i'}))}$$

- Take a set of N choices and let return to choice i be

$$U_i = \delta_i + \varepsilon_i$$

and choice be

$$U^* = \max\{U_1, U_2, ...U_N\}$$

Then the probability of choice i is

$$P_i = \frac{\exp(\delta_i)}{\sum_{i'=1}^{N} \exp(\delta_{i'})}$$

- Independence of irrelevant alternatives. Relative probability do i instead of i' is (independent of existence of other alternatives)

$$\frac{P_i}{P_{i'}} = \frac{\exp(\delta_i)}{\exp(\delta_{i'})}$$

- Formula for maximum utility is

$$U^* = \max\{U_1, U_2, ...U_N\} = \gamma + \log\left(\sum_{i=1}^{N} \exp(\delta_i)\right) \qquad (1)$$

where $\gamma \approx .5772$ is Euler's constant.

- Expected value of ε_i (unconditioned) is

$$E[\varepsilon_i] = \gamma$$

- Next, calculate expected value of ε_i, given choice i.

$$E[\varepsilon_i | a_i = 1] = \gamma - \log(P_i) \tag{2}$$

Back to the Replacement problem

- Use vector notation. Then write current state as (ω, ε). Let (ω_1,ε_1) be the initial state (period 1). Also will leave dependence on parameter vector θ implicit for now.

$$V(\omega_1, \varepsilon_1) = \max_{\{a_1,a_2,a_3,\ldots\}} E\left[\sum_{t=1}^{\infty} \beta^{t-1} \left(\pi_{\omega,a_t} + \varepsilon_{a_t,t}\right) |\omega_1, \varepsilon_{0,1}, \varepsilon_{1,1}\right]$$

- To construct the Bellman equation, define the *choice specific value function*

$$\begin{aligned}\tilde{V}(\omega, \varepsilon, a) &= \pi_{\omega,a} + \beta E_{\omega',\varepsilon'}[V(\omega', \varepsilon')|\omega, a] + \varepsilon_a \\ &= \pi_{\omega,a} + \beta E_{\omega'}[E_{\varepsilon}[V\left(\omega', \varepsilon')|\omega'\right]|\omega, a\big)] + \varepsilon_a\end{aligned}$$

So

$$V(\omega, \varepsilon) = \max\left\{\tilde{V}(\omega, \varepsilon, 0), \tilde{V}(\omega, \varepsilon, 1)\right\}$$

- Solve this problem recursively. We begin with a $V^{\circ}(\cdot,\cdot,\cdot)$ and use the Bellman equation to map to a new function $V'(\cdot,\cdot,\cdot)$. We need a starting value and a good way to get this is to temporarily assume there is a finite horizon and calculate the return in the terminal period. Calculating a starting value this way yields

$$V^{\circ}(\omega,\varepsilon) = \max\left\{\pi_{\omega,0} + \varepsilon_0, \pi_{\omega,1} + \varepsilon_1\right\}$$

Using result (1) specified above

$$V^{\circ}(\omega) \equiv E_{\varepsilon}\left[V(\omega,\varepsilon)\right] = \left[\gamma + \log\left(\exp(\pi_{\omega,0}) + \exp(\pi_{\omega,1})\right)\right.$$

the probability of replacement is

$$P^{\circ}_{1|\omega} = \frac{\exp(\pi_{\omega,1})}{\exp(\pi_{\omega,0}) + \exp(\pi_{\omega,1})} \tag{3}$$

- Now take arbitrary $V^{\circ}(\omega, \varepsilon)$, define

$$V'(\omega, \varepsilon) = \max\left\{\tilde{V}'(\omega, \varepsilon, 0), \tilde{V}'(\omega, \varepsilon, 1)\right\},$$

for

$$\begin{aligned} \tilde{V}'(\omega, \varepsilon, a) &= \pi_{\omega,a} + \beta E_{\omega'}[E_{\varepsilon}[V^{\circ}\left(\omega', \varepsilon'\right)|\omega']|\omega, a\Big)] + \varepsilon_a \\ &= \pi_{\omega,a} + \beta E_{\omega'}[V^{\circ}(\omega')|\omega, a] + \varepsilon_a \end{aligned}$$

- Iterate until convergence.

Estimation: Nested-Fixed Point Approach

- All of the above interative procedure is for a GIVEN θ. But how estimate θ?

- Estimate transition $\Pr(\omega'|\omega, a)$ directly from the observed transitions in the data. Given assumption above, no selection issues to worry about. (Would be an issue if there is measurement error on ω). Let's say we have $\Pr(\omega'|\omega, a)$ in hand, and turn our focus to estimating $\pi_{\omega,a}$.

- Nested fix point. Given θ, get fixed point of value function iteration to get solve for $V(\omega, \varepsilon, \boldsymbol{\theta})$. This gives of the choice specific value function for choice a.

$$\tilde{V}(\omega, \varepsilon, a, \theta) = \pi_{\omega,a} + \beta E_{\omega'}[E_{\varepsilon}[V\left(\omega', \varepsilon', \theta)|\omega'\right]|\omega, a\Big)] + \varepsilon_a$$

and we can calculate the probability of choice a, given ω,

$$P_{a|\omega}(\theta) = \frac{\exp(\pi_{\omega,a} + E_{\omega'}[E_{\varepsilon}[V(\omega', \varepsilon', \theta)|\omega']|\omega, a)])}{\sum_{b=0}^{1} \exp(\pi_{\omega,b} + E_{\omega'}[E_{\varepsilon}[V(\omega', \varepsilon', \theta)|\omega']|\omega, b)])}$$

- Take data on choice of a given ω and maximize the likelihood.

- More generally, let $\hat{P}_{a,\omega}$ be an estimate of the conditional probability of choice a given ω. Use some metric to pick θ so that $P_{a|\omega}(\theta)$ is close to $\hat{P}_{a,\omega}$.

- Issue about the "curse of dimensionality"

- Has led to “two-step” approaches that avoid the inter loop.

Two-Step Approaches

- Will start by going over the single agent problem above and use the two-step method. This approach is due to Hotz-Miller (1993).

- Want to say up front that the payoff from the two-step rather than the nest-fixed-point really comes in big when we go to oligopolistic interaction.

  – That is where the curse of dimensionality bites hard. (If A choices and N firms then A^N possible outcomes.

  – Two-step approach is an end-run around the multiplicity of equilibria issue which bits hard in oligpoly models. (Basically irrelevant in single-agent problems).

- Big idea. Start from the conditional choice probabilities (CCP) estimated from the data in a first step.

  - Now find parameters such that predicted behavior is consistent with the observed behavior.

  - Key advantage in oligopoly context is never have to calculate the equilibrium even once!

    * Agent is playing against other agents. Agent 1 needs to make predictions about how Agent 2 behaves given the state. How does that happen in the data? Can use this when studying agent 1's problem. So convert the entire analysis to single agent decision theory.

    * Of course this logic only works if when there a multiple equilibria, only a single one is being played in the data.

β

How Two Step Procedure Works

Step 1: Estimate the CCP $\hat{P}_{a,\omega}$, call the entire matrix $\widehat{CCP}$

with multiple agents $\hat{P}_{a,\omega,i}$

Step 2:

- Given ω at $t = 1$, and choice is a, calculate discounted fraction of time at state ω' in future periods. Note since we have $\hat{P}_{a,\omega}$, this will also give us the discounted fraction of time we are at ω' and choice is a'.
  - Can do this by simulation (as in BBL), may be easiest. Let $G_{\omega'|\omega}$ be discounted fraction of time at (ω') in future, given at ω now.
  - Example of one simulated path. Let s index a particular simulation. (S total number of simulations)

  - ∗ Start at ω. Then use $\hat{P}_{a,\omega}$ to draw a_1^s, then use $P\left(\omega'|\omega,a\right)$ to draw ω_2^s, then $\hat{P}_{a,\omega}$ to draw a_2.

  - ∗ Now take ω_t^s and a_t^s and iterate this to get ω_{t+1}^s and a_{t+1}^s. Stop after T periods

- – Now have $\omega_t^s(\omega)$ for S simulations and $2 \leq t \leq T$. Define indicator function $1_{[event]} = 1$ if event realized. Define

$$G_{\omega'|\omega} = \sum_{t=2}^{T} \beta^{t-1} \sum_{s=1}^{S} \frac{1_{[\omega_t^s(\omega)=\omega']}}{S}$$

- Note $G_{\omega''|\omega}$ doesn't depend upon θ, so can do this once.

- Define the choice specific value function not including unobservated shock (leaving it out, so we can plug it into that logit probability formulas)

$$\begin{aligned}\tilde{U}(\omega, a, \theta, \widehat{CCP}) &= \pi_{\omega,a} \\ &\quad + \sum_{\omega'=0}^{\bar{\omega}} \sum_{a'=0}^{1} \Big(\pi_{\omega',a'} + \gamma - \log(\hat{P}_{a',\omega'}\Big)) G_{\omega'|\omega}(\widehat{CCP}) \hat{P}_{a',\omega'}\end{aligned}$$

- Next observe we have a mapping in the space of CCP

$$\hat{P}_{a,\omega} = \frac{\exp(\tilde{U}(\omega, a, \theta, \widehat{CCP}))}{\exp(\tilde{U}(\omega, 0, \theta, \widehat{CCP}) + \exp(\tilde{U}(\omega, 1, \theta, \widehat{CCP}))}$$

- Pick θ to get this equation to hold as well as possibile.