1 Outline.

1. Motivation.

2. Estimation of Single Agent Dynamic Models.

3. Brute force

4. Two step methods

2 Estimation of Games.

- In the next module, we will discuss two-step estimation of dynamic games.

- Auguerreguberia and Mira(2007), Pesendorfer and Schmidt-Dengler(2006), Pakes, Ovstrovsky and Berry (2006), Bajari, Benkard and Levin (2007).

- It is now possible to estimate games that it would be difficult to compute even once.

- Games use a two step estimator.

- In the first step, the economist flexibly estimates, the probability distribution of the endogenous actions.

- If you are in equilibrium, since agents have rational expectations, we can consistently estimate agents' actual beliefs in many problems.

- In the second set, given these beliefs, we apply revealed preference to recover the utility parameters that make the observed actions optimal.

- These estimators are computationally light and simple to implement.

3 Overview of Rust.

- We will start off by considering single agent dynamic models.

- Reference, Rust, Handbook of Econometrics Volume IV.

- In standard random utility models, preference are a function of covariates, parameters and choice specific shocks.

- In dynamic discrete choice, one assumes that the payoffs are the discount sum of static logit models.

- Examples: Educational choice, fertility, dynamic demand.

4 The Model.

- The model is of a single agent making an optimal decision.

- Let $t = 1, ..., T, T \leq \infty$ denote time.

- The set S is the state space.

- The agent's choice set (or decision space) is D.

- The constraint set is $D(s_t) \subseteq D$.

- Transition probabilities (law of motion), $p(s_{t+1}|s_t, d_t)$, evolution of state variables.

- β, discount factor.

- Decision rule, map δ from $s \rightarrow D(s)$

- The period utility function is $u(s_t, d_t)$

- Given a decision rule δ, we can calculate expected discounted utility as:

$$EU(\delta) = E\left\{\sum \beta^t u(s_t, \delta(s_t))\right\}$$

- Agents choose δ in order to maximize expected discounted utility.

- Remark: The assumption that optimal decision rules are Markovian is without loss of generality.

- The proof is similar to SLP.

- Consider a decision problem with a large but finite value of T.

- At the last time period, the optimal decision only depends on the state because utility is additively separable.

- By induction, assume that at time period t that decision rules only depend on the state.

- Then at t-1, take the decision rule in future periods as given (this is the induction hypothesis).

- Since utility is additively separable, the optimal decision rule only depends on the current state.

- It is worth noting that this model makes some strong assumptions.

- People act as if they are able to solve a dynamic program.

- The evidence is mixed about how well we actually do this.

- This model also assumes that agents maximize expected utility with additively separable utility functions.

- Non-separable preferences, non-expected utility frameworks; robust decision theory; time inconsistent preferences.

5 Estimation: Brute Force Approaches.

- In the vast majority of applications, authors assume that utility takes the form:

$$u(s, d) = u(x, d) + \varepsilon(d)$$

- That is, we interpret the state as having two parts: $s = (x, \varepsilon)$.

- The variables x are the state variables observed to the economist.

- The variables $\varepsilon(d)$ are choice specific shocks (i.e. use the logit model).

- Let $\pi(y; x, d)$ denote the chain governing the evolution of the endogenous state variable.

- Rust encourages us to interpret $\varepsilon(d)$ as an unobserved state variables.

- This is slightly cheesy: few if any sensible, primitive economic models would generate this structure.

- Why would the unobserved state be an iid shock to specific choices?

- Rust notes that we can write the value function as follows:

$$\begin{aligned} V(x,\varepsilon) &= \max_{d \in D(s)} [\varepsilon(d) + v(x,d)] \\ v(x,d) &= u(x,d) + \beta \int V(y,\varepsilon)\pi(y;x,d)dF(\varepsilon) \end{aligned}$$

- We sometimes refer to $v(x, d)$ as the choice specific value function.

- It represents the continuation value from choosing d today and your optimal actions into the future.

- If we use the logit model, then we can turn finding $v(x, d)$ into a fixed point problem and apply the contraction mapping theorem:

$$\Psi(v) = u(x, d) + \beta \int \log \left(\sum_{d'} \exp v(y, d') \right) \pi(y; x, d) dF(\varepsilon)$$

- Given the choice specific value functions, the logit model implies

$$P(d|x) = \frac{\text{exp}(v(x,d))}{\sum_{d'} \text{exp}(v(x,d'))}$$

- Estimation is done using a likelihood based approach.

- Utility depends on parameters θ, $u(x,d;\theta)$.

- First you solve for the choice specific value functions given θ.

- Normally, we estimate law of motion $\pi(y;x,d)$ in a first stage.

- We then use value function iteration to solve for $v(x,d;\theta)$

$$\Psi(v(\theta)) = u(x, d; \theta) +$$
$$\beta \int \log \left(\sum_{d'} \exp v(y, d'; \theta) \right) \pi(y; x, d) dF(\varepsilon)$$

- Second, maximize the likelihood function in the "outer loop"

- The probability that d is chosen given x and θ is:

$$P(d|x) = \frac{\exp(v(x, d; \theta))}{\sum_{d'} \exp(v(x, d'; \theta))}$$

- Given data on x_t, d_t the MLE is found by minimizing:

$$L(\theta) = \frac{1}{T}\sum_t \log\left(\frac{\exp(v(x_t, d_t; \theta))}{\sum_{d'} \exp(v(x_t, d'; \theta))}\right)$$

- Some issues:

1. Curse of dimensionality in dyamic programming

- Time consuming to solve for max

- Tight parameterizations can be required

2. Evaluation of $L(\theta)$ is approximate because of numerical error in solving for $v(x_t, d_t; \theta)$

- Maximization may be subtle

6 Estimation: Two-Step Approaches.

- Hotz and Miller had the following insight.

- Algorithms of the type Rust applies are a bear to compute.

- Using standard methods, it may be possible to non-parametrically estimate $P(d|x)$.

- This is just the probability that some discrete event occurs conditional on x.

- Given two decision, d and d', the following must hold:

$$\frac{P(d|x)}{P(d'|x)} = \frac{\exp(v(x,d))}{\exp(v(x,d'))}$$

$$\log(P(d|x)) - \log(P(d'|x)) = v(x,d) - v(x,d')$$

- From the above equation, non-parametric estimates $\widehat{P}(d|x)$ can be used to estimate $\widehat{v}(x, d)$.

- Let $\widehat{P}(d|x)$ and $\widehat{\pi}(y; x, d)$ denote estimates of the choices probabilities and the law of motion.

- Let $u(x, d; \theta)$ denote a parametric representation of the utility function.

- Given θ and these estimates, we can simulate the choice specific value functions quite easily:

$$\widetilde{v}(s, d; \theta) = u(s, d; \theta) + \\ E \sum \left(\beta^t u(s', d'; \theta) + \varepsilon(d') | \widehat{P}(d|x), \widehat{\pi}(y; x, d) \right)$$

- We then choose θ to minimize the distance between $\widetilde{v}(s, d; \theta)$ and $\widehat{v}(x, d)$.

- To summarize, the steps are:

1. Non-parametrically estimate $\widehat{P}(d|x)$ using standard techniques.

2. Estimate the choice specific value functions by running the regression: $\log(\widehat{P}(d|x)) - \log(\widehat{P}(d'|x)) = \widehat{v}(x,d) - \widehat{v}(x,d')$

3. Choose θ to minimize the distance between $\widetilde{v}(s,d;\theta)$ and $\widehat{v}(x,d)$.

- Note that this procedure has a much lower computational burden than Rust's MLE based approach.

- Since it has a lower burden, we can add many additional bells and whistles to the estimator.

- This estimator also has the advantage of being very clear about where the identification is coming from.

- The parameters are being chosen to rationalize the implied choice probabilities.

- There are, however disadvantages.

- Since it is non-parametric, the data requirements are high.

- There is also probably a loss of efficiency from not computing the full equilibrium.

- Also, estimating using Rust's method gives the researchers a lot of intuition about the model.

- One has the chance to solve it many times.

6.1 Identification.

- Rust notes that these models are subject to an identification problem.

- We could set $\beta = 0$ and rationalize the data through a clever choice of $u(x, d)$.

- This can be seen in the formula, for the static model:

$$\frac{P(d|x)}{P(d'|x)} = \frac{\text{exp}(u(x,d))}{\text{exp}(u(x,d'))}$$

- However, $\beta = 0$ is clearly an unreasonable assumption.

- Interest rates are positive (for instance), indicating that people are impatient.

- Economists have a fairly good sense of what plausible values are for β (0.92 to 0.97).

- We shall explore identification in greater detail below.

7 Outline.

1. Motivation of Pesendorfer and Schmidt-Dengler.

2. The Model.

3. Estimation

8 Motivation.

- Next, we turn to the problem of estimating dynamic games.

- These are dynamic extensions of the static games we saw last time.

- A key insight is related to our study of static games.

- We can substitute in first stage estimates of the actions of others.

- Therefore, we do not need to compute the full equilibrium during estimation.

9 The Model.

- Dynamic Markov Perfect Equilibrium.

- Strategies only depend on payoff relevant state variables.

- The utility inthe stage game is defined using the logit model.

- Time is discrete, $t = 1, 2, 3, \ldots \infty$.

- There are $\mathbf{N} = \{1, \ldots, N\}$ firms.

- A typical firm is labeled i.

- The number of firms is fixed and does not change over time.

- A firm is endowed with a state at time t, $s_i^t \in S_i$

- Note that the set of states in this game is discrete.

- Each firm gets to take an action $a_i \in A_i = \{0, ..., K\}$ at each time period.

- All firms take their actions simultaneously.

- In addition to the observed state variables, each agent has a vector of unobserved state variables (logit errors).

- The errors $\varepsilon_i(k)$ are the private information shock associated with each action k

- The error terms are of course observed by the agents before the actions are taken.

- The cardinality of the vector of actions $a^t = (a_1^t, ..., a_N^t)$ is $(K+1)^N$.

- The transition for the state variables is a function $g : A \times S \times S \to [0, 1]$.

- The number $g(a, s^t, s^{t+1})$ is the probability of making a transition from s^t to s^{t+1} given the vector of actions a.

- Obviously, to generate a well defined chain we must have:

$$\sum_{s'} g(a, s, s') = 1$$

- The payoff to firm i from action $a_i = k$ at the stage game can be written as:

$$\pi_i(a_i = k, a_{-i}, s) + \varepsilon_i(k)$$

- Firms maximize expected discounted utility:

$$E_0 \left\{ \sum \beta^t \left(\pi_i(a_i, a_{-i}, s) + \varepsilon_i(k) \right) \right\}$$

- Focus on Markov Perfect equilibrium.

- Let $P(a_j|s)$ denote the probability distribution over firm j's actions given the state s.

10 Estimation

- In our applied work, we will make $\pi_i(a_i, a_{-i}, s; \theta)$ a function of parameters θ

- Assume that $\widehat{g}(a, s, s')$ is estimated in a first stage

- Define $\pi_i(a_i, s)$ as:

$$\pi_i(a_i, s) = \sum_{a_{-i}} \pi_i(a_i, a_{-i}, s) P(a_{-i}|s) + \varepsilon_i(a_i)$$

- Suppose that we estimate $\widehat{P}(a_{-i}|s)$ in a first step.

- Substituting into the above:

$$\pi_i(a_i, s) = \sum_{a_{-i}} \pi_i(a_i, a_{-i}, s) \widehat{P}(a_{-i}|s) + \varepsilon_i(a_i)$$

- Let the trasition rule be:

$$f(a_i, s, s') = \sum_{a_{-i}} g(a_i, a_{-i} s, s') \widehat{P}(a_{-i}|s)$$

- We have turned our mulitple agent problem into a single agent problem.

- We could now do Hotz-Miller in order to estimate the dynamic game.

- Some important extensions:

1. Existence and Identification (Pesendorfer Schmidt-Dengler)

2. Efficiency (Aguirregabiria and Mira)

3. Unobserved Heterogeneity (Kasahara and Shimotsu)

11 Dynamic Games w/ continuous choices

- Next we consider dynamic games where the choice variables are continuous and/or discrete.

- Use forward simulation method of Bajari, Benkard and Levin.

Notation.

- Assume discrete state space and discrete action space (for convenience only).

- Agents: $i = 1, ..., N$

- Time: $t = 1, ..., \infty$

- States: $\mathbf{s}_t \in S \subset R^G$, commonly known.

- Actions: $a_{it} \in A_i$, simultaneously chosen.

- Transitions: $P(\mathbf{s}_{t+1}|\mathbf{a}_t, \mathbf{s}_t)$.

- Discount Factor: β (known to econometrician).

Objective Function: Agent maximizes EDV,

$$\mathbb{E}_t \sum_{t=0}^{\infty} \beta^t u_i(\mathbf{a}_t, \mathbf{s}_t). \tag{1}$$

Equilibrium.

- Concept: Markov Perfect Equilibrium [MPE]

- Strategies: $\sigma_i : S \rightarrow A_i$.

- Recursive Formulation:

  $V_i(\mathrm{s}|\boldsymbol{\sigma}) = u_i(\boldsymbol{\sigma}(\mathrm{s}), \mathrm{s}) + \beta \int V_i(\mathrm{s}'|\boldsymbol{\sigma}) dP(\mathrm{s}'|\sigma(\mathrm{s}), \mathrm{s})$

- A MPE is given by a Markov profile, σ, such that for all i, s, σ_i'

$$V_i(\mathrm{s}|\sigma_i, \boldsymbol{\sigma}_{-i}) \geq V_i(\mathrm{s}|\sigma_i', \boldsymbol{\sigma}_{-i})$$

First Step.

- Estimate policy functions,

$$\sigma_i : S \rightarrow A_i$$

- and state transition function,

$$P : S \times A \rightarrow \Delta(S).$$

- Often will also estimate “static” parts of period return.

Examples:

- Production functions, (Olley-Pakes)
- Investment policies, (nonparametric)
- Entry/Exit policies, (nonparametric)
- Labor Supply,
- Static supply-demand system (BLP)
- State transitions: (parametric/nonparametric).

Second Step.

- Idea: Find the set of parameters that rationalize the data.

- I.e., conditional on P and σ, find the set of parameters that satisfy the requirements for equilibrium.

- Optimality Inequalities: For all i, σ_i', and initial state, s_0, it must be that

$$\mathbb{E}_{\sigma_i,\boldsymbol{\sigma}_{-i}|\mathbf{s}_0} \sum_{t=0}^{\infty} \beta^t u_i(\mathbf{a}_t, \mathbf{s}_t) \geq \mathbb{E}_{\sigma_i',\boldsymbol{\sigma}_{-i}|\mathbf{s}_0} \sum_{t=0}^{\infty} \beta^t u_i(\mathbf{a}_t, \mathbf{s}_t), \tag{3}$$

- The system of inequalities, (3), contains all information available from the definition of equilibrium.

- Assume: period return function is linear in the parameters
(stronger than needed),

$$u_i(\mathbf{a}, \mathrm{s}; \boldsymbol{\theta}) = \boldsymbol{\Phi}_i(\mathbf{a}, \mathrm{s}) \cdot \boldsymbol{\theta}. \qquad (4)$$

- Let

$$W(\mathrm{s}_0; \widehat{\sigma}_i, \widehat{\boldsymbol{\sigma}}_{-i}) = \widehat{\mathbb{E}}_{\sigma_i, \boldsymbol{\sigma}_{-i} | \mathrm{s}_0} \sum_{t=0}^{\infty} \beta^t \boldsymbol{\Phi}_i(\mathbf{a}_t, \mathrm{s}_t).$$

- That is, given our policy functions, we simulate the integral $\widehat{\mathbb{E}}_{\sigma_i, \boldsymbol{\sigma}_{-i} | \mathrm{s}_0} \sum_{t=0}^{\infty} \beta^t \boldsymbol{\Phi}_i(\mathbf{a}_t, \mathrm{s}_t)$ by Monte Carlo along the lines discussed in chapter 12.

- Then the system (3) can be written as,

$$W(\mathrm{s}_0; \sigma_i, \boldsymbol{\sigma}_{-i}) \cdot \boldsymbol{\theta} \geq W(\mathrm{s}_0; \sigma_i', \boldsymbol{\sigma}_{-i}) \cdot \boldsymbol{\theta}, \qquad (5)$$

- for all i, σ_i', s_0.

- Given σ and P, let Θ_0 be the set of parameters that rationalize the observed data,

  $\Theta_0(\boldsymbol{\sigma},P) := \{\boldsymbol{\theta} : \boldsymbol{\theta}, \boldsymbol{\sigma}, P$ satisfy (5) for all $s_0, i, \sigma_i'\}$.

  where (5) is the system of optimality inequalities,

  $$W(\mathrm{s}_0; \sigma_i, \boldsymbol{\sigma}_{-i}) \cdot \boldsymbol{\theta} \geq W(\mathrm{s}_0; \sigma_i', \boldsymbol{\sigma}_{-i}) \cdot \boldsymbol{\theta}, \quad (5)$$

  for all i, σ_i', s_0.

- The goal of estimation is to learn Θ_0.

- We cover two cases:

1. Θ_0 is a singleton (model is identified).

2. Θ_0 is a set (model is not identified).

Identified Case.

- Let observed policy function be $\sigma = (\sigma_1, ..., \sigma_N)$.

- To simplify notation, abstract away from estimation error (easy to fix using standard theory of two-step estimators).

- Consider a finite set of alternative policies, σ_i' that agent i could have chosen, but did not.

- Define:

$$g(x, \theta) = \left[\widehat{W}(\mathrm{s}; \sigma_i, \boldsymbol{\sigma}_{-i}) - \widehat{W}(\mathrm{s}; \sigma_i', \boldsymbol{\sigma}_{-i})\right] \cdot \boldsymbol{\theta}$$

- Let n_I be the number of alternative (but not chosen) revealed preference inequalities that we consider.

- Abusing notation, let $g(X_k, \theta)$ denote the function g evaluated at a particular revealed preference inequality.

- Minimize:

$$Q_n(\theta) = \frac{1}{n} \sum_i 1 \{g(X_k, \theta) < 0\} \, g(X_k, \theta)^2$$

Comments:

- Two-step estimator is constructed by evaluating Q_n usiing first stage policy function estimates.

- Variances in denominator are known because they can be estimated very precisely using the simulation draws.

- Computationally light because simulation only needs to be done once, prior to maximizing likelihood.

- All second stage error comes from simulation.

Dynamic Oligopoly w/ Investment

- Like Pakes and McGuire or Ericson and Pakes.

- Demand:

$$U_{rj} = \gamma_0 \ln(z_i) + \gamma_1 \ln(y - p_i) + \varepsilon_{rj}$$

- Where z_i is product quality- integer valued.

- ε_{rj} is iid logit error term.

- Constant marginal costs, μ.

- Bertrand price competition- estimate this using standard demnd estimation.

- Investment $I_{it} \in R_+$ is successful w/ prob:

$$\frac{\rho I_{it}}{1 + \rho I_{it}}$$

- where ρ is a parameter.

- If investment is successful, product quality increases by 1.

- The cost of investment is $C(I) = \xi I$

- There is also an outside good who's quality will move up w/ prob δ each period.

- Scrap value for exit, ϕ.

- $\chi_i(s_t)$ is exit policy decision.

- One potential entrant each period.

- v_{et} is private entry cost drawn from distribution, $[v^L, v^H]$

- State variable is number of firms and vector of product qualities.

- The value function for an incumbent firm is:

$$
\begin{aligned}
V_i(s,\sigma) &= W^1(s,\sigma) + W^2(s,\sigma) + W^3(s,\sigma) \\
W^1(s,\sigma) &= E\left[\sum \beta^t \pi_i(s_t)|s\right] \\
W^2(s,\sigma) &= -E\left[\sum \beta^t I_i(s_t)|s\right]\xi \\
W^3(s,\sigma) &= -E\left[\sum \beta \chi_i(s_t)|s\right]\phi
\end{aligned}
$$

First Stage.

- Estimate demand parameters using MLE and markup using first order conditions for optimal pricing.

- Estimate the transition parameters, ρ and δ also by MLE.

- Estimate investment and exit policy functions using local linear regressions.

Second Stage.

- For every initial state, s_0, and every alternative investment policy, $\sigma'(\mathrm{s}) = (I'(\mathrm{s}), \chi'(\mathrm{s}))$,

$$\left[\widehat{\mathbb{E}}_{\sigma_i,\sigma_{-i}} \sum_{t=0}^{\infty} \beta^t \pi_i(s_t) - \widehat{\mathbb{E}}_{\sigma'_i,\sigma_{-i}} \sum_{t=0}^{\infty} \beta^t \pi_i(s_t)\right]$$

$$+\xi \cdot \left[\widehat{\mathbb{E}}_{\sigma_i,\sigma_{-i}} \sum_{t=0}^{\infty} \beta^t I_i(s_t) - \widehat{\mathbb{E}}_{\sigma'_i,\sigma_{-i}} \sum_{t=0}^{\infty} \beta^t I_i(s_t)\right]$$

$$+\phi \left[\begin{array}{c} \widehat{\mathbb{E}}_{\sigma_i,\sigma_{-i}} \sum_{t=0}^{\infty} \beta^t \{\chi(\mathrm{s}_t) = 1\} \\ -\widehat{\mathbb{E}}_{\sigma'_i,\sigma_{-i}} \sum_{t=0}^{\infty} \beta^t \{\chi'(\mathrm{s}_t) = 1\} \end{array}\right] \geqslant 0$$

- To get alternative policies, add mean zero error term to the investment and exit policies.

- Also straightforward to estimate sunk cost of entry distribution (parametrically or nonparametrically) – see paper for details.

- Advantages of approach:
  - Computationally light.
  - Simple to apply and uses standard techniques.
  - More generally applicable than existing methods.
- Disadvantages of approach:
  - Not efficient.
  - Data requirement?