Sequential Estimation of Dynamic Discrete Games

Victor Aguirregabiria and Pedro Mira

presented by

Yang, Yong Hyeon

DYNAMIC GAMES

- Variables
  - x_t: observable state variable
  - ε_t: unobservable state variable
  - a_t: observable decision variable (discrete, e.g. binary)
- Firms choose optimal a_{it} that maximizes

$$V_i(x_t, \varepsilon_{it}) = \max_{a_{it}} E\left[\beta^{s-t} \sum_{s=t}^{\infty} \widetilde{\Pi}_i(a_s, x_s, \varepsilon_{is}) \middle| x_t, \varepsilon_{it}\right]$$

where expectations are over others' strategies and ε_{jt}'s, meaning that

  - Firms' profit depends on other firms' characteristic and decision.
  - Firms observe their own private information before decision.
  - Firms' private information affects their profit only.

DYNAMIC PROGRAMMING

- Usually we denote

$$V_i(x_t, \varepsilon_{it}) = \max_{a_{it}} E\left[\widetilde{\Pi}_i(a_t, x_t, \varepsilon_{it}) + \beta V_i(x_{t+1}, \varepsilon_{it+1}) \middle| x_t, \varepsilon_{it}\right]$$

or simply drop time subscripts to write

$$V_i(x, \varepsilon_i) = \max_{a_i} E\left[\widetilde{\Pi}_i(a, x, \varepsilon_i) + \beta V_i(x', \varepsilon_i') \middle| x, \varepsilon_i\right]$$

- $\widetilde{\Pi}$ is known

- V is unique given $\widetilde{\Pi}$ and distriubition of x' and ε_i'.

ESTIMATION

- Firms' decision depends on parameters.
- Estimate parameters that make firms' decision seem reasonable!
  - GMM

$$E[a_i - P_i(x;\theta)|x] = 0$$

    meaning that given x, firm i chooses $a_i = 1$ with probability $P_i(x;\theta)$.
  - MLE

$$\max \ \log L(\varepsilon;\theta) = \log L(a,x;\theta)$$

- Basically this requires us to solve for V_i explicitly.
  - $P_i(x;\theta)$ can be derived from $V_i(x,\varepsilon_i)$.
  - $L(a,x;\theta)$ depends on likelihood of a, which can be obtained from $V_i(x,\varepsilon_i)$.

PROBLEM

- This task requires much time.
  - Need to numerically solve for V_i using contraction mapping.
  - Need to simulate integrals (or use quadrature).
  - Need to repeat these for many candidate parameters.

$$V_i(x, \varepsilon_i; \theta) = \max_{a_i} E\left[\widetilde{\Pi}_i(a, x, \varepsilon_i; \theta) + \beta V_i(x', \varepsilon_i') \middle| x, \varepsilon_i; \theta\right]$$

- Still doable, but want to reduce time cost.

SOLUTION

- Try to bypass the step for contraction mapping and get $P_i(x)$ or $L(a, x; \theta)$ directly using data.
  - Hotz and Miller (1993): GMM in the context of individual dynamic decision
  - Hotz, Miller, Sanders and Smith (1994): include simulation
  - Pakes, Ostrovsky, and Berry (2004): extend to dynamic games
  - Aguirregabiria and Mira (2002): MLE, individual dynamics, iteration
- Solve for V on some set of parameters and use it repeatedly
  - Ackerberg (2001): use importance sampling

SUMMARY OF THIS PAPER

- Extend Aguirregabiria and Mira (2002) to dynamic games
  - Need to consider strategic complexity
- Try to incorporate unobserved heterogeneity
  - Hotz and Miller cannot be applied to the case where unobserved heterogeneous variable exists.
    (to be explained later)
- Conclusion: Nested Pseudo Likelihood (NPL) estimator is useful when
  - multiple equilibria exist,
  - state spaces are large, and/or
  - conditional choice probabilities are not easily estimated
    (first and third will be explained later)

MODEL IN DETAIL

- Assume additive separability

$$\widetilde{\Pi}_i(a_t, x_t, \varepsilon_{it}) = \Pi_i(a_t, x_t) + \varepsilon_{it}(a_{it})$$

- Assume conditional independence

$$\begin{aligned} &p(x_{t+1}, \varepsilon_{t+1}|a_t, x_t, \varepsilon_t) \\ &= p(\varepsilon_{t+1}|x_{t+1}, a_t, x_t, \varepsilon_t)p(x_{t+1}|a_t, x_t, \varepsilon_t) \\ &= p_\varepsilon(\varepsilon_{t+1})f(x_{t+1}|a_t, x_t) \end{aligned}$$

- Assume independent error across players

- Assume discrete and finite x_t

VALUE FUNCTION

- Original value function is

$$V_i(x, \varepsilon_i) = \max_{a_i} E\left[\Pi_i(a, x) + \varepsilon_i(a_i) + \beta V_i(x', \varepsilon_i') \middle| x, \varepsilon_i\right]$$

- Let $\sigma = \{\sigma_i\}_{i=1}^N$ be strategy profiles players are using.

$$V_i^\sigma(x, \varepsilon_i) = \max_{a_i} E\left[\Pi_i^\sigma(a, x) + \varepsilon_i(a_i) + \beta V_i^\sigma(x', \varepsilon_i') \middle| x, \varepsilon_i\right]$$

- Let $\pi_i^\sigma(a_i, x) = E[\Pi_i^\sigma(a, x)]$ be the expected profit.

$$V_i^\sigma(x, \varepsilon_i) = \max_{a_i} \left\{\pi_i^\sigma(a_i, x) + \varepsilon_i(a_i) + \beta E\left[V_i^\sigma(x', \varepsilon_i') | x, \varepsilon_i\right]\right\}$$

- Explicitly writing integrals,

$$V_i^\sigma(x, \varepsilon_i) = \max_{a_i} \left\{\pi_i^\sigma(a_i, x) + \varepsilon_i(a_i) \right.$$
$$\left. + \beta \sum_{x'} f^\sigma(x' | a_i, x) \int V_i^\sigma(x', \varepsilon_i') g_i(\varepsilon_i') d\varepsilon_i' \right\}$$

CONDITIONAL CHOICE PROBABILITY

- Define $v_i^\sigma(a_i, x)$ as

$$v_i^\sigma(a_i, x) = \pi_i^\sigma(a_i, x) + \beta \sum_{x'} f^\sigma(x'|a_i, x) \int V_i^\sigma(x', \varepsilon_i')g_i(\varepsilon_i')d\varepsilon_i'$$

so that

$$V_i^\sigma(x, \varepsilon_i) = \max_{a_i} \left\{ v_i^\sigma(a_i, x) + \varepsilon_i(a_i) \right\}$$

- Conditional choice probability $P_i(x)$ is defined as

$$P_i(x) = E[a_i|x] = \Pr(a_i = 1|x) = \Pr\left(v_i^\sigma(1, x) + \varepsilon_i(1) > v_i^\sigma(0, x) + \varepsilon_i(0) \right)$$

INTEGRATED VALUE FUNCTION

- Use integrated value function to get $P_i(x)$.

$$V_i^\sigma(x) = \int V_i^\sigma(x, \varepsilon_i) g_i(\varepsilon_i) d\varepsilon_i$$
$$= \int \max_{a_i} \Big\{ v_i^\sigma(a_i, x) + \varepsilon_i(a_i) \Big\} g_i(\varepsilon_i) d\varepsilon_i$$

where now

$$v_i^\sigma(a_i, x) = \pi_i^\sigma(a_i, x) + \beta \sum_{x'} f^\sigma(x'|a_i, x) V_i^\sigma(x')$$

- Why is $V_i^\sigma(x)$ useful?
  - Decrease in the number of state spaces
  - Easy to calculate $v_i^\sigma(a_i, x)$: do not need to do integral
    $\longrightarrow$ Easy to calculate $P_i(x)$: e.g. Type I extreme value distribution implies

$$P_i(x) = \frac{e^{v_i^\sigma(1,x)}}{e^{v_i^\sigma(0,x)} + e^{v_i^\sigma(1,x)}}$$

METHODOLOGY

- Need to do contraction mapping to get $V_i^\sigma(x)$?

$$V_i^\sigma(x) = \int \max_{a_i} \left\{ \pi_i^\sigma(a_i, x) + \varepsilon_i(a_i) + \beta \sum_{x'} f^\sigma(x'|a_i, x) V_i^\sigma(x') \right\} g_i(\varepsilon_i) d\varepsilon_i$$

- NO. Then we would have the same computational burden.
- Alternatively, if we have an estimate of $v_i^\sigma(a_i, x)$, say, $\widehat{v_i^\sigma}(a_i, x)$,

$$\widehat{V_i^\sigma}(x) = \int \max_{a_i} \left\{ \widehat{v_i^\sigma}(a_i, x) + \varepsilon_i(a_i) \right\} g_i(\varepsilon_i) d\varepsilon_i$$

  is a good estimate of $V_i^\sigma(x)$.
- Unfortunately, we do not have an estimate of $v_i^\sigma(a_i, x)$.
- But we can estimate $P_i(x)$. Can we use this?
- YES. It is actually a part of strategies σ.

ALTERNATIVE WAY

- Find a way to go from $P_i(x)$ to $V_i^\sigma(x)$.

$$V_i^\sigma(x) = \int \max_{a_i} \left\{ \pi_i^\sigma(a_i, x) + \varepsilon_i(a_i) + \beta \sum_{x'} f^\sigma(x'|a_i, x) V_i^\sigma(x') \right\} g_i(\varepsilon_i) d\varepsilon_i$$

$$= P_i(x) \left[\pi_i^\sigma(1, x) + E[\varepsilon_i(1)|a_i = 1, x] + \beta \sum_{x'} f^\sigma(x'|a_i = 1, x) V_i^\sigma(x') \right]$$

$$+ [1 - P_i(x)] \left[\pi_i^\sigma(0, x) + E[\varepsilon_i(0)|a_i = 0, x] + \beta \sum_{x'} f^\sigma(x'|a_i = 0, x) V_i^\sigma(x') \right]$$

- Use estimated conditional choice probability $\widehat{P}_i(x)$ to get

$$\widehat{V_i^P}(x) = \widehat{P}_i(x) \left[\widehat{\pi_i^P}(1, x) + E[\varepsilon_i(1)|a_i = 1, x] + \beta \sum_{x'} \widehat{f^P}(x'|a_i = 1, x) \widehat{V_i^P}(x') \right]$$

$$+ [1 - \widehat{P}_i(x)] \left[\widehat{\pi_i^P}(0, x) + E[\varepsilon_i(0)|a_i = 0, x] + \beta \sum_{x'} \widehat{f^P}(x'|a_i = 0, x) \widehat{V_i^P}(x') \right]$$

- Since x is finite, $\widehat{V_i^P}(x)$ can be easily obtained using matrix calculation.

SUMMARIZE THE PROCEDURE

(1) Estimate $\widehat{P}_i(x)$.

(2) Obtain $\widehat{\pi_i^P}(a_i, x; \theta)$ and $\widehat{f^P}(x'|a_i, x; \theta)$

(3) Solve for $\widehat{V_i^P}(x; \theta)$

This is a good estimate of $V_i^\sigma(x; \theta)$.

(4) Compute $\widehat{v_i^P}(a_i, x; \theta)$

$$\widehat{v_i^P}(a_i, x; \theta) = \widehat{\pi_i^P}(a_i, x; \theta) + \beta \sum_{x'} \widehat{f^P}(x'|a_i, x; \theta) \widehat{V_i^P}(x'; \theta)$$

(5) Derive $\widetilde{P}_i(x; \theta)$: e.g. Type I extreme value distribution implies

$$\widetilde{P}_i(x; \theta) = \frac{e^{\widehat{v_i^P}(1,x;\theta)}}{e^{\widehat{v_i^P}(0,x;\theta)} + e^{\widehat{v_i^P}(1,x;\theta)}}$$

(6) (Hotz and Miller) Estimate parameters that satisfies

$$E\left[a_i - \widetilde{P}_i(x; \theta) \middle| x\right] = 0$$

THEORETICAL REPRESENTATION

- A Markov Perfect Equilibrium as a Fixed Point
  - Let Ψ be a best response mapping.
  - Then, P, a set of conditional choice probabilities, is a MPE iff

$$P = \Psi(P)$$

    meaning a fixed point of Ψ.

- Put true P, then we would get P in the last under true parameters.

INTRODUCTION TO ESTIMATION

- Hotz and Miller, and Pakes, Ostrovsky and Berry use GMM.
- Aguirregabiria and Mira use MLE.
  - But this is not a standard MLE.
  - Maximize likelihood of discrete (e.g. binary) choice

$$\max_{\theta} \sum_{m=1}^{M} \sum_{t=1}^{T} \sum_{i=1}^{N} \log \widetilde{P}_i(x_{mt})^{a_{imt}} \left[1 - \widetilde{P}_i(x_{mt})\right]^{1-a_{imt}}$$

    where m denotes markets.
  - Since $\widetilde{P}_i(x)$ is not a true probability when θ is not a true parameter, we call this the pseudo likelihood function, and the maximizer is called the pseudo maximum likelihood (PML) estimator.

Note: The precise definition of pseudo likelihood function is broader

$$Q_M(\theta, P) = \frac{1}{M} \sum_{m=1}^{M} \sum_{t=1}^{T} \sum_{i=1}^{N} \log \Psi_i(a_{imt} | x_{imt}; P, \theta)$$

where Ψ_i is a best response mapping given the initial P.

PML ESTIMATION

- Infeasible PML estimation
  - Suppose we know true P. Use this to get $\widetilde{P}$.
  - Form a pseudo likelihood function, and find θ maximizing it.

$$\sqrt{M}(\widehat{\theta} - \theta) \xrightarrow{d} N(0, \Omega_{\theta\theta}^{-1})$$

    where $\Omega_{\theta\theta} \equiv E[\nabla_\theta s_m \nabla_\theta s_m']$ is the information matrix.

- 2 step PML estimation
  - Estimate $\widehat{P}$ nonparametrically so that $\sqrt{M}(\widehat{P} - P) \xrightarrow{d} N(0, \Sigma)$.
  - Use this to get $\widetilde{P}$.
  - Form a pseudo likelihood function, and find θ maximizing it.

$$\sqrt{M}(\widehat{\theta} - \theta) \xrightarrow{d} N(0, V_{2S})$$

    where $V_{2S} = \Omega_{\theta\theta}^{-1} + \Omega_{\theta\theta}^{-1}\Omega_{\theta P}\Sigma\Omega_{\theta P}'\Omega_{\theta\theta}^{-1}$
  - Identification assumptions are required.

IDENTIFICATION ASSUMPTION

- Denote true P by P^0.
- Assume

(A) For every observation (m, t), P^0 is played.

(B) Players expect P^0 to be played in future periods.

(C) For any $\theta \neq \theta^0$, the solution of $P = \Psi(P)$ is $P \neq P^0$.

(D) The observations are independent across markets.

DRAWBACK OF THE PML ESTIMATOR

- Inefficiency
  - V_{2S} depends on Σ.
  - If $\widehat{P}$ is inefficient, $\widehat{\theta}$ is inefficient, too.
- Bias
  - Imprecise $\widehat{P}$ makes huge bias, especially in small samples.
  - e.g. if we use crude frequency (accept/reject) estimator.
- Feasibility
  - $\widehat{P}$ may not be estimated.
  - e.g. if unobserved heterogeneous variable exists.
- Motivates another estimator.

FIXED POINT OF THE PROCEDURE

- Iterations on Procedure
  - Start with $\widehat{P}_0$.
  - Use this to get $\widetilde{P}(x;\theta)$. Find $\widehat{\theta}_1$ maximizing the PL function.
  - Define $\widehat{P}_1(x) = \widetilde{P}(x;\widehat{\theta}_1)$.
  - Use this to get $\widetilde{P}(x;\theta)$. Find $\widehat{\theta}_2$ maximizing the PL function. $\cdots$
  - Define $\widehat{P}_{K-1}(x) = \widetilde{P}(x;\widehat{\theta}_{K-1})$.
  - Use this to get $\widetilde{P}(x;\theta)$. Find $\widehat{\theta}_K$ maximizing the PL function.
- Includes 2 step PML ($K = 1$ with $\widehat{P}_0 = \widehat{P}$)
- As $K \to \infty$, the sequence $\{\widehat{\theta}_K, \widehat{P}_K\}$ converges.
  - Put its limit, then we get the same: a fixed point of the procedure.
  - If its limit is singleton, $\widehat{\theta}_\infty$ is consistent.

NESTED PML ESTIMATION

- Nested PML (NPL) estimator is defined as follows.
  - Find all NPL fixed points.
  - Find the maximizer of PL function among them. Call it the NPL estimator.
- Properties
  - Consistent: $\widehat{\theta} \xrightarrow{p} \theta$
  - Asymptotic Normality

$$\sqrt{M}(\widehat{\theta} - \theta) \xrightarrow{d} N(0, V_{NPL})$$

  - More efficient than the infeasible PML estimator under some condition.
  - Less finite sample bias.
  - Any initial $\widehat{P}_0$ leads to the NPL estimator $\widehat{\theta}$.
    So can be applied to the case where $\widehat{P}$ may not be estimated.
  - Existence of multiple equilibria does not lead to multiple NPL fixed points.
  - Still require the same identification assumptions.

MONTE CARLO EXPERIMENT

- Use the model of firms' entry and exit.
  - 5 firms choose to enter or exit upon 160 states.
  - Try to estimate cost and revenue parameters.
- Compare infeasible/crude frequency 2 step/logit 2 step PML estimators with the NPL estimator.
  - NPL always converges (faster with logit initial estimator).
  - It always converges to the same estimate regardless of initial guess.
  - Crude frequency 2 step PML has a very large bias.
  - NPL has sometimes less variance and MSE than infeasible PML.
  - Logit 2 step PML is pretty good.
- Using a smooth nonparametric estimator in the 1st step reduces bias in the 2nd step. (Pakes, Ostrovsky, and Berry)

PERMANENT UNOBSERVED HETEROGENEITY

- What if a state variable is unobserved and heterogeneous but permanent?
  - Let ω_m is a time-invariant market characteristic.
  - Common knowledge to the players but unobserved to the econometrician.
  - Need to construct $P_i(x, \omega)$ to use the same method.
  - Problem: cannot construct $P_i(x, \omega)$ since ω is not observed.
- Still doable with the original value function method.
- Infeasible with Hotz and Miller style estimation.
- Aguirregabiria and Mira propose the method and the condition.

CONDITION AND METHOD

- Assume conditional iid
  - Let $\overline{x}_m$ be time invariant market characteristics.
  - Given $\overline{x}_m$, ω_m is iid across markets: $\Pr(\omega_m = \omega^l | \overline{x}_m) = \varphi_l(\overline{x}_m)$
  - ω_m does not affect $f(x'|a, x)$
- The idea is that
  - write ω_m as a function of observable state variables and parameters.
  - residual is iid, so we can do ML type estimation.

$$\begin{aligned}
\log \Pr(\text{data}) &= \sum_{m=1}^{M} \log \Pr(a_m, x_m) \\
&= \sum_{m=1}^{M} \log \left(\sum_{l=1}^{L} \varphi_l(\overline{x}_m) \Pr(a_m, x_m | \omega^l) \right)
\end{aligned}$$

EMPIRICAL APPLICATION

- Use Chilean industries of five sectors.
  - Firms choose to enter or exit
  - There is a permanent unobserved heterogeneity ω_m taking on 21 values.
- Results
  - NPL always converges.
  - It always converges to the same estimate.
  - Unreasonable estimate without ω_m: it says that a firm's profit increases with the number of firms.
  - Including ω_m makes reasonable estimation.
  - Actual interpretation of parameters is not interesting.