ECO 2901
EMPIRICAL INDUSTRIAL ORGANIZATION

Lecture 9:
Structure and Estimation of
Dynamic Games of Oligopoly Competition

Victor Aguirregabiria (University of Toronto)

March 14, 2019

Lecture 9: Structure and estimation of dynamic games of oligopoly competition

1. Introduction
2. Structure of empirical dynamic games
3. Data, Identification, and Estimation
4. Counterfactuals with Multiple Equilibria

1. Introduction

Dynamic Games: Introduction

- In oligopoly industries, firms compete in investment decisions that:
  - have returns in the future (forward-looking);
  - involve substantial uncertainty;
  - have important effects on competitors 'profits (competition / game)
- Some examples are:
  - Investment in R&D, innovation.
  - Investment in capacity, physical capital.
  - Product design / quality
  - Market entry / exit ...

Dynamic Games: Introduction [2]

• Measuring and understanding the **dynamic strategic interactions** between firms decisions (e.g., dynamic complementarity or substitutability) is important to understand the forces behind the dynamics of an industry or to evaluate policies.

• Investment costs, uncertainty, and competition effects play an important role in these decisions.

• Structural estimation of these parameters is necessary for some empirical questions.

• Empirical dynamic games provide a framework to estimate these parameters and perform policy analysis.

Examples of Empirical Applications

• Competition in R&D and product innovation between Intel and AMD: Goettler and Gordon (JPE, 2011).

• Product innovation: incumbents & new entrants (hard drive industry): Igami (JPE, 2017).

• Land use regulation and entry-exit in the hotel industry: Suzuki (IER; 2013).

• Environmental regulation, entry-exit and capacity in cement industry: Ryan (ECMA, 2012).

• Subsidies to entry in small markets of the dentist industry: Dunne et al. (RAND, 2013);

Examples of Empirical Applications [2]

• Fees for musical performance & choice of format of radio stations:
Sweeting (ECMA, 2013).

• Hub-and-spoke networks and entry-exit in the airline industry:
Aguirregabiria and Ho (JoE, 2012).

• Dynamic price competition:
Kano (IJIO, 2013); Ellickson, Misra, and Nair (JMR, 2012).

• Cannibalization and preemption strategies in fast-food industry:
Igami and Yang (QE, 2016).

• Demand uncertainty and firm investment in the concrete industry:
Collard-Wexler (ECMA, 2013);

Examples of Empirical Applications [3]

- Release date of a movie: Einav (EI, 2010).

- Time-to-build, investment, and uncertainty in the shipping industry: Kalouptsidi (AER, 2014).

- Endogenous mergers: Jeziorski (RAND, 2014).

- Exploitation of a common natural resource (fishing): Huang and Smith (AER, 2014).

2. Structure of Dynamic Games of Oligopoly Competition

Dynamic Games: Basic Structure

- Time is discrete and indexed by t.

- The game is played by N firms that we index by i.

- Following the standard structure in the **Ericson-Pakes (1995) framework**, firms compete in two different dimensions: a static dimension and a dynamic dimension.

- For instance: given the state of the industry at period t firms compete in prices (static competition), and decide the quality of their products (dynamic investment decision).

Dynamic Games: Basic Structure (2)

• The investment decision can be an entry/exit decision, a choice of capacity, investment in equipment, R&D, product quality, other product characteristics, etc.

• As an example, I use here a model of competition in product quality that is similar to Pakes & Mcguire (RAND, 1994).

• The action is taken to maximize the expected and discounted flow of profits in the market,

$$E_t \left(\sum_{s=0}^{\infty} \delta^s \; \Pi_{it+s} \right)$$

where $\delta \in (0,1)$ is the discount factor, and Π_{it} is firm i's profit at period t.

Decision variable

• Let a_{it} be the variable that represents the investment decision of firm i at period t.

• As an example, I use here a model of competition in product quality that is similar to Pakes & Mcguire (RAND, 1996).

• $a_{it} \in \{0, 1, ..., A\}$,
$a_{it} = 0$: firm i is not active in the market;
$a_{it} = a > 0$: firm i is active with a product of quality a.

• Consumers willingness to pay for the product of firm i, v_i, and the cost of product i (marginal and fixed) depend on the quality levels:

$$v_i(1) < v_i(2) < ... < v_i(A)$$
$$c_i(1) < c_i(2) < ... < c_i(A)$$

State variables

• At every period t the industry can be described in terms of three sets of state variables affecting firms' profits:

$$\mathbf{a}_{t-1},\ \mathbf{z}_t,\ \boldsymbol{\varepsilon}_t$$

• **Endogenous common knowledge state variables:**
$\mathbf{a}_{t-1} = (a_{1t-1}, a_{2t-1}, ..., a_{Nt-1})$. Vector with firms' product qualities at previous period. They affect profit because there is a cost of changing/adjusting the level of quality.

• **Exogenous common knowledge state variables:** $\mathbf{z}_t$, affecting demand and costs.

• **Exogenous private common knowledge state variables:**
$\boldsymbol{\varepsilon}_t = (\varepsilon_{1t}, \varepsilon_{2t}, ..., \varepsilon_{Nt})$, affecting firms' costs. ε_{it} is private info of firm i.

Profit function

- The profits of firm i at time t are given by

$$\Pi_{it} = VP_{it} - FC_{it} - IC_{it}$$

where:

VP_{it} represents variable profit;
FC_{it} is the fixed cost of operating;
IC_{it} is an investment / adjustment cost

Variable profit function

- The variable profit VP_{it} is comes from the equilibrium of a static Bertrand game:

$$VP_{it} = (p_{it} - c_i(z_t)) \; q_{it}$$

p_{it} and q_{it} are the price and the quantity sold by firm i.

- According this model, the quantity is:

$$q_{it} = H_t \; \frac{1\{a_{it} > 0\} \; \exp\{v_i(a_{it}, z_t) - \alpha \; p_{it}\}}{1 + \sum_{j=1}^{N} 1\{a_{jt} > 0\} \; \exp\{v_j(a_{jt}, z_t) - \alpha \; p_{jt}\}}$$

where H_t is the number of consumers in the market (market size).

- Bertrand equilibrium implies the "indirect" variable profit function:

$$\theta_i^{VP}(\mathbf{a}_t, \mathbf{z}_t) = (p_i^*[\mathbf{a}_t, \mathbf{z}_t] - c_i(z_t)) \; q_i^*[\mathbf{a}_t, \mathbf{z}_t]$$

Fixed cost

- The fixed cost is paid every period that the firm is active in the market:

$$FC_{it} = \theta_i^{FC}(a_{it}, \mathbf{z}_t) + \varepsilon_{it}^{FC}(a_{it})$$

- $\theta_i^{FC}(a, \mathbf{z}_t)$ is the fixed cost of firm i if the quality of its product is a.

- $\varepsilon_{it}^{FC}(a)$ are zero-mean shocks that are private information of firm i.

Fixed cost (2)

• There are two main reasons why we incorporate private information shocks in the model.

• [1] As shown in Doraszelski and Satterthwaite (2012), it is a way to guarantee that the dynamic game has at least one equilibrium in pure strategies.

• [2] They are convenient econometric errors. If private information shocks are independent over time and over players, and unobserved to the researcher, they can 'explain' players heterogeneous behavior without generating endogeneity problems.

Investment / Adjustment costs

- There are costs of adjusting the level of quality:

$$IC_{it} = \theta_i^{AC}(a_{it} - a_{it-1},\ a_{it-1}; \mathbf{z}_t) + \varepsilon_{it}^{AC}(a_{it})$$

- $\theta_i^{AC}(a_{it} - a_{it-1},\ a_{it-1}; \mathbf{z}_t)$ is the adjustment cost function, such that:

  $\theta_i^{AC}(0, a_{it-1}; \mathbf{z}_t) = 0$
  $\theta_i^{AC}(\Delta, a_{it-1}; \mathbf{z}_t) > 0$ if $\Delta \neq 0$.
  If $a_{it-1} = 0$, this adjustment cost is actually the cost of market entry;
  If $a_{it-1} > 0$ and $a_{it} = 0$, this adjustment cost is actually the cost of market exit.

- $\varepsilon_{it}^{AC}(a)$ is a private information shock in the investment cost

Profit function

- In summary, the profit function has the following structure:

$$\Pi_{it} = \pi_i\left(a_{it}, \mathbf{a}_{-it}, a_{it-1}, \mathbf{z}_t\right) - \varepsilon_{it}(a_{it})$$

where:

$$\pi_i\left(a_{it}, \mathbf{a}_{-it}, a_{it-1}, \mathbf{z}_t\right) = \theta_i^{VP}(a_{it}, \mathbf{a}_{-it}, \mathbf{z}_t) - \theta_i^{FC}(a_{it}, \mathbf{z}_t) - \theta_i^{AC}(a_{it} - a_{it-}$$

and:

$$\varepsilon_{it}(a_{it}) = \varepsilon_{it}^{FC}(a_{it}) + \varepsilon_{it}^{AC}(a_{it})$$

Evolution of the state variables

• (1) **Exogenous state variables:** follow an exogenous Markov process with transition probability function $F_z(\mathbf{z}_{t+1}|\mathbf{z}_t)$.

• (2) **Endogeneous state variables:** In this example, this transition is deterministic and very simple. Choices at period t (a_{it}) determined (are equal to) the endogenous state variables at $t+1$ (a_{it}).

• (3) **Private information state variables**. ε_{it} is i.i.d. over time and independent across firms with CDF G_i.

Timing of decisions and state variables

• In this example, I consider that firms' decisions about the quality of their products are made at the beginning of period t and they are already effective at period t such that VP_{it} depends on $\mathbf{a}_t$.

• An alternative timing that has been considered in many applications is that there is a **one-period time-to-build**. The decision is made at period t, and entry costs are paid at period t, but the quality choice at period t is not effective until period $t+1$. This is in fact the timing of decisions in Ericson and Pakes (1995), or Pakes & McGuire (1994).

• All the results below can be easily generalized to this model with time-to-build.

Markov Perfect Equilibrium

• Most of the recent literature in IO studying industry dynamics focuses on studying a Markov Perfect Equilibrium (MPE), as defined by Maskin and Tirole (Econometrica, 1988).

• The key assumption in this solution concept is that players' strategies are functions of only payoff-relevant state variables.

• In this model, the payoff-relevant state variables for firm i are $(\mathbf{a}_{t-1}, \mathbf{z}_t, \varepsilon_{it})$.

• We use $\mathbf{x}_t$ to represent the vector of common knowledge state variables:

$$\mathbf{x}_t \equiv (\mathbf{a}_{t-1}, \mathbf{z}_t)$$

Markov Perfect Equilibrium (2)

• Let $\boldsymbol{\alpha} = \{\alpha_i(\mathbf{x}_t, \varepsilon_{it}) : i \in \{1, 2, ..., N\}\}$ be a set of strategy functions, one for each firm.

• A MPE is an N-tuple of strategy functions $\boldsymbol{\alpha}$ such that every firm is maximizing its value given the strategies of the other players.

• For given strategies of the other firms, the decision problem of a firm is a single-agent dynamic programming (DP) problem.

Markov Perfect Equilibrium (3)

• Let $V_i^{\boldsymbol{\alpha}}(\mathbf{x}_t, \varepsilon_{it})$ be the value function of the DP problem that describes the best response of firm i to the strategies $\boldsymbol{\alpha}_{-i}$ of the other firms.

• This value function is the unique solution to the Bellman equation:

$$V_i^{\boldsymbol{\alpha}}(\mathbf{x}_t, \varepsilon_{it}) = \max_{a_{it}} \left\{ \begin{array}{l} \Pi_i^{\boldsymbol{\alpha}}(a_{it}, \mathbf{x}_t) - \varepsilon_{it}(a_{it}) \\ \\ +\delta \displaystyle\int V_i^{\boldsymbol{\alpha}}(\mathbf{x}_{t+1}, \varepsilon_{it+1})\ dG_i(\varepsilon_{it+1})\ F_i^{\boldsymbol{\alpha}}(\mathbf{x}_{t+1} | a_{it}, \mathbf{x}_t) \end{array} \right\}$$

$\Pi_i^{\boldsymbol{\alpha}}(a_{it}, \mathbf{x}_t) =$ One-period profit for firm i given the strategies of the other firms.
$F_i^{\boldsymbol{\alpha}}(\mathbf{x}_{t+1} | a_{it}, \mathbf{x}_t) =$ Transition probability of state variables given the strategies of the other firms.

Markov Perfect Equilibrium (4)

- The expected one-period profit $\Pi_i^{\boldsymbol{\alpha}}(a_{it}, \mathbf{x}_t)$ is:

$$\Pi_i^{\boldsymbol{\alpha}}(a_{it}, \mathbf{x}_t) \quad = \quad \sum_{\mathbf{a}_{-it}} \left[\prod_{j \neq i} \Pr\left(\alpha_j(\mathbf{x}_t, \varepsilon_{jt}) = a_{jt} \mid \mathbf{x}_t \right) \right] \pi_i \left(a_{it}, \mathbf{a}_{-it}, \mathbf{x}_t \right)$$

And the expected transition of the state variables is:

$$F_i^{\boldsymbol{\alpha}}(\mathbf{x}_{t+1} | a_{it}, \mathbf{x}_t) \quad = \quad F_z(\mathbf{z}_{t+1} | \mathbf{z}_t) \prod_{j \neq i} \Pr\left(\alpha_j(\mathbf{x}_t, \varepsilon_{jt}) = a_{jt} \mid \mathbf{x}_t \right)$$

Markov Perfect Equilibrium (5)

• A Markov perfect equilibrium (MPE) in this game is an N-tuple of strategy functions $\boldsymbol{\alpha}$ such that for any player i and for any $(\mathbf{x}_t, \varepsilon_{it})$ we have that:

$$\alpha_i(\mathbf{x}_t, \varepsilon_{it}) = \arg\max_{a_{it}} \{v_i^{\boldsymbol{\alpha}}(a_{it}, \mathbf{x}_t) - \varepsilon_{it}(a_{it})\}$$

with

$$v_i^{\boldsymbol{\alpha}}(a_{it}, \mathbf{x}_t) \equiv \Pi_i^{\boldsymbol{\alpha}}(a_{it}, \mathbf{x}_t) + \delta \int V_i^{\boldsymbol{\alpha}}(\mathbf{x}_{t+1}, \varepsilon_{it+1}) \; dG_i(\varepsilon_{it+1}) \; F_i^{\boldsymbol{\alpha}}(\mathbf{x}_{t+1} | a_{it}, \mathbf{x}_t)$$

Conditional Choice Probabilities

- Given a strategy function $\alpha_i(\mathbf{x}_t, \varepsilon_{it})$, we can define the corresponding *Conditional Choice Probability (CCP)* function as :

$$\begin{aligned} P_i(a|\mathbf{x}) &\equiv \Pr\left(\alpha_i(\mathbf{x}_t, \varepsilon_{it}) = a \mid \mathbf{x}_t = \mathbf{x}\right) \\ &= \int 1\{\alpha_i(\mathbf{x}_t, \varepsilon_{it}) = a\} \; dG_i(\varepsilon_{it}) \end{aligned}$$

Conditional Choice Probabilities (2)

• Since choice probabilities are integrated over the continuous variables in ε_{it}, they are lower dimensional objects than the strategies $\boldsymbol{\alpha}$.

• For instance, when both a_{it} and $\mathbf{x}_t$ are discrete, CCPs can be described as vectors in a finite dimensional Euclidean space.

• There is a one-to-one relationship between strategy functions $\alpha_i(\mathbf{x}_t, \varepsilon_{it})$ and CCP functions $P_i(a|\mathbf{x}_t)$.

• We can use $\Pi_i^{\mathbf{P}}$ and $F_i^{\mathbf{P}}$ instead of $\Pi_i^{\boldsymbol{\alpha}}$ and $F_i^{\boldsymbol{\alpha}}$ to represent the expected profit function and the transition probability function, respectively.

MPE in terms of CCPs

- A MPE is a vector of CCPs, $\mathbf{P} \equiv \{P_i(a|\mathbf{x}) : \text{for any } (i, a, x)\}$, such that:

$$P_i(a|\mathbf{x}) = \Pr\left(a = \arg\max_{a_i}\left\{v_i^{\mathbf{P}}(a_i, \mathbf{x}) - \varepsilon_i(a_i)\right\} \mid \mathbf{x}\right)$$

- $v_i^{\mathbf{P}}(a_i, \mathbf{x})$ is a conditional choice probability function, but it has a slightly different definition that before. Now, $v_i^{\mathbf{P}}(a_i, \mathbf{x})$ represents the value of firm i if the firm chooses alternative a_i today and

all the firms, including firm i, behave according to their respective CCPs in $\mathbf{P}$.

- The Representation Lemma in Aguirregabiria and Mira (2007) shows that every MPE in this dynamic game can be represented using this mapping.

MPE in terms of CCPs (2)

- The form of this equilibrium mapping depends on the distribution of ε_i.

- For instance, in the entry/exit model, if ε_i is $N(0,1)$:

$$P_i(1|\mathbf{x}) = \Phi\left(v_i^{\mathbf{P}}(1,\mathbf{x}) - v_i^{\mathbf{P}}(0,\mathbf{x})\right)$$

- In the model with endogenous quality choice, if $\varepsilon_i(a)$'s are extreme value type 1 distributed:

$$P_i(a|\mathbf{x}) = \frac{\exp\left\{v_i^{\mathbf{P}}(a,\mathbf{x})\right\}}{\sum_{a'=0}^{A}\exp\left\{v_i^{\mathbf{P}}(a',\mathbf{x})\right\}}$$

Computing values and best response probs

- By definition:

$$v_i^{\mathbf{P}}(a_i, \mathbf{x}) = \Pi_i^{\mathbf{P}}(a_i, \mathbf{x}) + \delta \sum_{\mathbf{x}'} V_i^{\mathbf{P}}(\mathbf{x}') \ F_i^{\mathbf{P}}(\mathbf{x}'|a_i, \mathbf{x})$$

- $V_i^{\mathbf{P}}(\mathbf{x})$ is the value of firm i if all the firms, including firm i, behave according to their CCPs in $\mathbf{P}$, and the current state is $\mathbf{x}$.

- By definition, $V_i^{\mathbf{P}}$ is the unique solution of the recursive expression:

$$V_i^{\mathbf{P}}(\mathbf{x}) = \sum_{a_i=0}^{A} P_i(a_i|\mathbf{x}) \left[\Pi_i^{\mathbf{P}}(a_i, \mathbf{x}) + \delta \sum_{\mathbf{x}'} V_i^{\mathbf{P}}(\mathbf{x}') \ F^{\mathbf{P}}(\mathbf{x}'|a_i, \mathbf{x})\right]$$

Computing values and best response probs [2]

- When the space $\mathcal{X}$ is discrete we can write this problem in vector form:

$$\mathbf{V}_i^{\mathbf{P}} = \sum_{a_i=0}^{A} \mathbf{P}_i(a_i) * \left[\text{˝}_i^{\mathbf{P}}(a_i) + \delta \, \mathbf{F}_i^{\mathbf{P}}(a_i) \, \mathbf{V}_i^{\mathbf{P}} \right]$$

$\mathbf{V}_i^{\mathbf{P}}$: is a $|\mathcal{X}| \times 1$ vector of values;
$\mathbf{P}_i(a_i)$: is a $|\mathcal{X}| \times 1$ vector of CCPs;
$\text{˝}_i^{\mathbf{P}}(a_i)$: is a $|\mathcal{X}| \times 1$ vector of expected payoffs;
$\mathbf{F}_i^{\mathbf{P}}(a_i)$: is a $|\mathcal{X}| \times |\mathcal{X}|$ matrix of transition probabilities.
$*$ is the Hadamard or element-by-element product

Computing values and best response probs [3]

$$\mathbf{V}_i^{\mathbf{P}} = \sum_{a_i=0}^{A} \mathbf{P}_i(a_i) * \left[''^{\mathbf{P}}_i(a_i) + \delta \, \mathbf{F}_i^{\mathbf{P}}(a_i) \, \mathbf{V}_i^{\mathbf{P}} \right]$$

- This is a linear system and the solution is:

$$\mathbf{V}_i^{\mathbf{P}} = \left(\mathbf{I} - \delta \left[\sum_{a_i=0}^{A} \mathbf{P}_i(a_i) * \mathbf{F}_i^{\mathbf{P}}(a_i) \right] \right)^{-1} \left(\sum_{a_i=0}^{A} \mathbf{P}_i(a_i) * ''^{\mathbf{P}}_i(a_i) \right)$$

Computing values and best response probs [4]

- Suppose that:

$$\pi_i \left(a_{it}, \mathbf{a}_{-it}, \mathbf{x}_t \right) = h_i \left(a_{it}, \mathbf{a}_{-it}, \mathbf{x}_t \right) \ \boldsymbol{\theta}_i$$

where $h_i \left(a_{it}, \mathbf{a}_{-it}, \mathbf{x}_t \right)$ is a vector that is known to the researcher; and $\boldsymbol{\theta}_i$ is a vector of structural parameters.

- The best response probabilities are:

$$P_i(a|\mathbf{x}_t) = \Pr \left(a = \arg\max_{a_{it}} \left\{ \widetilde{h}_i^{\mathbf{P}}(a_{it}, \mathbf{x}_t) \ \boldsymbol{\theta}_i - \varepsilon_{it}(a_{it}) \right\} \ | \ \mathbf{x}_t \right)$$

with $\widetilde{h}_i^{\mathbf{P}}(a_i, \mathbf{x})$ is the expected discounted value of the current and future

stream of values $h_i \left(a_{it+s}, \mathbf{a}_{-it+s}, \mathbf{x}_{t+s} \right)$.

- $\widetilde{h}_i^{\mathbf{P}}(a_i, \mathbf{x})$ has a closed form expression using the solution to the linear system presented above.

Computing values and best response probs [5]

• Therefore, with $\varepsilon_i(a)$'s extreme value type 1 distributed:

$$P_i(a|\mathbf{x}_t) = \frac{\exp\left\{\widetilde{h}_i^{\mathbf{P}}(a, \mathbf{x}_t)\ \boldsymbol{\theta}_i\right\}}{\sum_{a'=0}^{A} \exp\left\{\widetilde{h}_i^{\mathbf{P}}(a', \mathbf{x}_t)\ \boldsymbol{\theta}_i\right\}}$$

• Note that this best response probabilities define the equilibrium CCPs as the solution to a fixed point problem in the space of the vector of CCPs.

• Given this representation, it is clear by Brower's theorem than a MPE exists.

• This representation will be useful for the identification and estimation of the model.

3. Data, Identification, Estimation

Data

• The researcher observes a random sample of M markets, indexed by m, over T periods of time, where the observed variables consists of players' actions and state variables.

• For the moment, we consider that the industry and the data are such that:

(a) each firm is observed making decisions in every of the M markets;

(b) the researcher knows all the payoff relevant market characteristics that are common knowledge to the firms, $\mathbf{x}$.

Data

• With this type of data we can allow for rich firm heterogeneity that is fixed across markets and time by estimating firm-specific structural parameters, $\boldsymbol{\theta}_i$.

• This 'fixed-effect' approach to deal with firm heterogeneity is not feasible in data sets where most of the competitors can be characterized as *local players*, i.e., firms specialized in operating in a few markets.

• Condition (b) rules out the existence of unobserved market heterogeneity. Though it is a convenient assumption, it is also unrealistic for most applications in empirical IO. Later I present estimation methods that relax conditions (a) and (b) and deal with unobserved market and firm heterogeneity.

Data

• Suppose that we have a random sample of M local markets, indexed by m, over T periods of time, where we observe:

$$Data = \{\mathbf{a}_{mt},\ \mathbf{x}_{mt} : m = 1, 2, ..., M;\ t = 1, 2, ..., T\}$$

• We want to use these data to estimate the model parameters in the population that has generated this data: $\boldsymbol{\theta}^0 = \{\boldsymbol{\theta}_i^0 : i \in I\}$.

Identification

• A significant part of this literature has considered the following identification assumptions.

Assumption (ID 1): Single equilibrium in the data. Every observation in the sample comes from the same Markov Perfect Equilibrium, i.e., for any observation (m, t), $\mathbf{P}_{mt}^{0} = \mathbf{P}^{0}$.

Assumption (ID 2): No unobserved common-knowledge variables. The only unobservables for the econometrician are the private information shocks ε_{imt} and the structural parameters $\boldsymbol{\theta}$.

Identification (2)

- First, let's summarize the structure of the dynamic game of oligopoly competition.

- Let θ be the vector of structural parameters of the model, where $\theta = \{\theta_i : i = 1, 2, ..., N\}$.

- Let $\mathbf{P}(\theta) = \{P_i(a|\mathbf{x}, \theta) : \text{for any } (i, a, \mathbf{x})\}$ be a MPE of the model associated with θ. $\mathbf{P}(\theta)$ is a solution to the following equilibrium mapping: for any $(i, a_i, \mathbf{x})$:

$$P_i(a_i|\mathbf{x},\theta) = \frac{\exp\left\{\widetilde{h}_i^{\mathbf{P}}(a_i, \mathbf{x})\ \boldsymbol{\theta}_i\right\}}{\sum_{a'=0}^{A} \exp\left\{\widetilde{h}_i^{\mathbf{P}}(a', \mathbf{x})\ \boldsymbol{\theta}_i\right\}}$$

Identification (3)

• Under assumptions ID-1 & ID-2, the equilibrium that has generated the data, $\mathbf{P}^0$, can be estimated consistently and nonparametrically from the data. For any $(i, a_i, \mathbf{x})$:

$$P_i^0(a_i|\mathbf{x}) = \Pr(a_{imt} = a_i \mid \mathbf{x}_{mt} = \mathbf{x})$$

For instance, we can estimate consistently $P_i^0(a_i|\mathbf{x})$ using the following simple kernel estimator:

$$P_i^0(a_i|\mathbf{x}) = \frac{\sum_{m,t} 1\{a_{imt} = a_i\} \; K\left(\dfrac{\mathbf{x}_{mt} - \mathbf{x}}{b_n}\right)}{\sum_{m,t} K\left(\dfrac{\mathbf{x}_{mt} - \mathbf{x}}{b_n}\right)}$$

Identification (4)

- Second, given that $\mathbf{P}^0$ is identified, we can identify also the expected present values $\widetilde{h}_i^{\mathbf{P}^0}(a_i, \mathbf{x})$ at the "true" equilibrium in the population.

- Third, we know that $\mathbf{P}^0$ is an equilibrium associated to $\boldsymbol{\theta}^0$. Therefore, the following equilibrium conditions should hold: for any $(i, a_i, \mathbf{x})$,

$$P_i^0(a_i|\mathbf{x}) = \frac{\exp\left\{\widetilde{h}_i^{\mathbf{P}^0}(a_i, \mathbf{x})\ \boldsymbol{\theta}_i^0\right\}}{\sum_{a'=0}^{A} \exp\left\{\widetilde{h}_i^{\mathbf{P}^0}(a', \mathbf{x})\ \boldsymbol{\theta}_i^0\right\}}$$

- These equilibrium conditions identify $\boldsymbol{\theta}_i^0$.

Identification (5)

- For instance, in this logit example, we have that for $(i, a_i, \mathbf{x})$,

$$\ln\left(\frac{P_i^0(a_i|\mathbf{x})}{P_i^0(0|\mathbf{x})}\right) = \left[\widetilde{h}_i^{\mathbf{P}^0}(a_i, \mathbf{x}) - \widetilde{h}_i^{\mathbf{P}^0}(0, \mathbf{x})\right] \boldsymbol{\theta}_i^0$$

- Define $\mathbf{Y}_i \equiv$ vector with $\ln\left(\frac{P_i^0(a_i|\mathbf{x})}{P_i^0(0|\mathbf{x})}\right)$ for every value of $(a_i, \mathbf{x})$; $\mathbf{H}_i \equiv$ matrix with $\widetilde{h}_i^{\mathbf{P}^0}(a_i, \mathbf{x}) - \widetilde{h}_i^{\mathbf{P}^0}(0, \mathbf{x})$ for every value of $(a_i, \mathbf{x})$. Then,

$$\mathbf{Y}_i = \mathbf{H}_i \ \boldsymbol{\theta}_i^0$$

- Matrix $\mathbf{H}_i$ is full-column-rank such that $\boldsymbol{\theta}_i^0$ is identified as:

$$\boldsymbol{\theta}_i^0 = \left[\mathbf{H}_i' \ \mathbf{H}_i\right]^{-1} \left[\mathbf{H}_i' \ \mathbf{Y}_i\right]$$

Estimation

- We consider the following estimators:

  1. Two-step estimator

  2. Recursive K-step and NPL

  3. Simulation-Based estimation to approximate present values

Estimation (2)

- For the sake of concreteness, we consider the binary choice entry-exit game, where $\varepsilon_{it}(1) - \varepsilon_{it}(0)$ is iid $N(0,1)$.
- The equilibrium mapping is:

$$P_i(1|\mathbf{x}, \boldsymbol{\theta}_i) = \Phi\left(\left[\widetilde{h}_i^{\mathbf{P}}(1, \mathbf{x}) - \widetilde{h}_i^{\mathbf{P}}(0, \mathbf{x})\right] \boldsymbol{\theta}_i\right)$$

Pseudo Likelihood Function

- For the description of the different estimators, it is convenient to define the following **Pseudo Likelihood function**:

$$
\begin{aligned}
Q(\boldsymbol{\theta}, \mathbf{P}) \quad = \quad & \sum_{m=1}^{M} \sum_{i=1}^{N} \sum_{t=1}^{T} a_{imt} \; \ln \Phi \left(\left[\widetilde{h}_i^{\mathbf{P}}(1, \mathbf{x}_{mt}) - \widetilde{h}_i^{\mathbf{P}}(0, \mathbf{x}_{mt}) \right] \; \boldsymbol{\theta}_i \right) \\
& + (1 - a_{imt}) \ln \left[1 - \Phi \left(\left[\widetilde{h}_i^{\mathbf{P}}(1, \mathbf{x}_{mt}) - \widetilde{h}_i^{\mathbf{P}}(0, \mathbf{x}_{mt}) \right] \; \boldsymbol{\theta}_i \right) \right]
\end{aligned}
$$

- This pseudo likelihood function treats firms' beliefs $\mathbf{P}$ as parameters to estimate together with $\boldsymbol{\theta}$.
- Note that for given $\mathbf{P}$, the function $Q(\boldsymbol{\theta}, \mathbf{P})$ is the likelihood of a Probit model.

Two-step methods

- Suppose that we knew the equilibrium in the population, $\mathbf{P}^0$.
- Given $\mathbf{P}^0$ we can construct the variables $\widetilde{h}_i^{\mathbf{P}}(1, \mathbf{x}_{mt}) - \widetilde{h}_i^{\mathbf{P}}(0, \mathbf{x}_{mt})$ and then obtain a very simple estimator of $\boldsymbol{\theta}^0$.

$$\hat{\boldsymbol{\theta}} = \arg\max_{\boldsymbol{\theta}} \; Q(\boldsymbol{\theta}, \mathbf{P}^0)$$

- This estimator is root-M consistent and asymptotically normal under the standard regularity conditions. It is not efficient because it does not impose the equilibrium constraints (only asymptotically).
- While equilibrium probabilities are not unique functions of structural parameters, the best response probabilities that appear in $Q(\boldsymbol{\theta}, \mathbf{P})$ are unique functions of structural parameters and players' beliefs.

Two-step methods (2)

- The previous method is infeasible because $\mathbf{P}^0$ is unknown.
- However, under the Assumptions **"No-unobserved-market-heterogeneity"** and **"One-MPE-in-the-data"** we can estimate $\mathbf{P}^0$ consistently and at with a convergence rate such that the two-step estimator $\hat{\theta}$ is root-M consistent and asymptotically normal.
- For instance, a kernel estimator of $\mathbf{P}^0$ is:

$$\widehat{P}_i^0(\mathbf{x}) = \frac{\sum_{m=1}^{M} \sum_{t=1}^{T} a_{imt} \; K\left(\dfrac{\mathbf{x}_{mt} - \mathbf{x}}{b}\right)}{\sum_{m=1}^{M} \sum_{t=1}^{T} K\left(\dfrac{\mathbf{x}_{mt} - \mathbf{x}}{b}\right)}$$

Two-step methods: Finite sample properties (1)

- The most attractive feature of two-step methods is their relative simplicity.
- However, they suffer of a potentially important problem of finite sample bias.
- The finite sample bias of the two-step estimator of θ^0 depends very importantly on the properties of the first-step estimator of $\mathbf{P}^0$. In particular, it depends on the rate of convergence and on the variance and bias of $\widehat{\mathbf{P}^0}$.
- It is well-known that there is a **curse of dimensionality in the NP estimation** of a regression function such as $\mathbf{P}^0$.

Recursive K-step estimator

- K-step extension of the 2-step estimator. Given an initial consistent (NP) estimator $\widehat{\mathbf{P}^0}$, the sequence of estimators $\{\widehat{\boldsymbol{\theta}^K}, \widehat{\mathbf{P}^K} : K \geq 1\}$ is defined as:

$$\widehat{\boldsymbol{\theta}^{K+1}} = \arg\max_{\boldsymbol{\theta}} \; Q\left(\boldsymbol{\theta}, \widehat{\mathbf{P}^K}\right)$$

where:

$$\widehat{P_i^K}(\mathbf{x}) = \Phi\left(\left[\widetilde{h}_i^{\widehat{\mathbf{P}^{K-1}}}(1, \mathbf{x}) - \widetilde{h}_i^{\widehat{\mathbf{P}^{K-1}}}(0, \mathbf{x})\right] \; \widehat{\boldsymbol{\theta}_i^K}\right)$$

- Aguirregabiria and Mira (2002, 2007) present Monte Carlo experiments which illustrate how this recursive estimators can have significantly smaller bias than the two-step estimator.
- Kasahara and Shimotsu (2008) derive a second order approximation to the bias of these K-stage estimators. They show that, if the equilibrium in the population is stable, then this recursive procedure reduces the bias.

Simulation-Based Estimation (1)

- Though two-step methods (with either PML or MI) are computationally much cheaper than full solution-estimation methods, they are still impractical for applications where the dimension of the state space X is very large, e.g., a discrete state space with millions of points or a model in which some of the observable state variables are continuous.
- To deal with this problem, Hotz, Miller, Sanders and Smith (REStud, 1994) proposed an estimator that uses simulation techniques to approximate the values $\widetilde{h}_{imt}^{\mathbf{P}}(a_i)$.
- In the context of dynamic games, Bajari, Benkard and Levin (BBL) have proposed to used this simulation and have extended it to models with continuous decision variables.

Simulation-Based Estimation (2)

- [For notational simplicity, omit subindexes i, m]. Remember that:

$$\widetilde{h}_t^{\mathbf{P}}(a_t) \equiv h_t^{\mathbf{P}}(a_t) + E\left(\sum_{s=1}^{\infty} \beta^s \; h_{t+s}^{\mathbf{P}}(a_{t+s}) \mid \mathbf{x}_t, \; a_t\right)$$

- The expectations $E(.)$ are taken over all the possible future paths of actions and state variables conditional on $(\mathbf{x}_t, a_t)$ and conditional on future behavior $\mathbf{P}$.

- The simulator of $\widetilde{h}_t^{\mathbf{P}}(a_t)$ ia obtained by replacing the true expectations $E(.)$ by a Monte Carlo approximation.

Simulation-Based Estimation (3)

- Starting at $(\mathbf{x}_t, a_t)$, we use the probabilities in $\mathbf{P}$, and the transition probabilities in F, to generate R simulated paths of future actions and state variables from period $t+1$ to $t+T^*$ (i.e., T^* periods ahead).

- We index simulated paths by $r \in \{1, 2, ..., R\}$. The $r-th$ simulated path is
$$\{a_{t+s}^{(r)}, \ \mathbf{x}_{t+s}^{(r)} : s = 1, 2, ..., T^*\}$$

Simulation-Based Estimation (4)

- A simulated path $\{a_{t+s}^{(r)}, \mathbf{x}_{t+s}^{(r)} : s = 1, 2, ..., T^*\}$ is obtained as follows.
- Given $(a_t, \mathbf{x}_t)$, we use the transition probability function $F^P(.|a_t, \mathbf{x}_t)$ to obtain a random draw $\mathbf{x}_{t+1}^{(r)}$.
- Given $\mathbf{x}_{t+1}^{(r)}$, we use the choice probability $P(\mathbf{x}_{t+1}^{(r)})$ to obtain a random draw $a_{t+1}^{(r)}$.
- Given $(a_{t+1}^{(r)}, \mathbf{x}_{t+1}^{(r)})$, we use the transition probability function $F^P(.|a_{t+1}^{(r)}, \mathbf{x}_{t+1}^{(r)})$ to obtain a random draw $\mathbf{x}_{t+2}^{(r)}$.
- And so on.

Simulation-Based Estimation (5)

- Then, given the simulated paths $\{a_{t+s}^{(r)}, \mathbf{x}_{t+s}^{(r)} : s = 1, 2, ..., T^*\}$, we construct the simulator of $\widetilde{h}_t^{\mathbf{P}}(a_t)$ as:

$$\widetilde{h}_t^{\mathbf{P},\mathbf{sim}}(a_t) \quad = \quad h_t^{\mathbf{P}}(a_t) + \frac{1}{R} \sum_{r=1}^{R} \left[\sum_{j=1}^{T^*} \delta^j \ h^{\mathbf{P}}(a_{t+s}^{(r)}, \mathbf{x}_{t+s}^{(r)}) \right]$$

- If the DP problem has finite horizon, or if T^* is large enough such that the approximation error associated with the truncation of paths is negligible, then these simulators are unbiased.

4. Counterfactual experiments with multiple equilibria

Counterfactual Experiments (1)

- One of the most attractive features of structural models is that they can be used to predict the effects of new policies or changes in parameters (counterfactuals).

- However, this a challenging exercise in a model with multiple equilibria.

- The data can identify the "factual" equilibrium. However, under the counterfactual scenario, which of the multiple equilibria we should choose?

Counterfactual Experiments (2)

- Different approaches have been implemented in practice.
- Select the equilibrium to which we converge by iterating in the (counterfactual) equilibrium mapping starting with the factual equilibrium $\mathbf{P}^0$
- Select the equilibrium with maximum total profits (or alternatively, with maximum welfare).
- Homotopy method: Aguirregabiria and Ho (2007)

Counterfactual Experiments: Homotopy method

- Let $\boldsymbol{\theta}$ be the vector of structural parameters in the model. An let $\Psi(\boldsymbol{\theta}, \mathbf{P})$ be the equilibrium mapping such that an equilibrium associated with $\boldsymbol{\theta}$ can be represented as a fixed point:

$$\mathbf{P} = \Psi(\boldsymbol{\theta}, \mathbf{P})$$

- The model could be completed with an equilibrium selection mechanism: i.e., a criterion that selects one and only one equilibrium for each possible $\boldsymbol{\theta}$.
- Suppose that there is a "true" equilibrium selection mechanism in the population under study, but we do not know that mechanism.
- Our approach here (both for the estimation and for counterfactual experiments) is completely agnostic with respect to the equilibrium selection mechanism.

Counterfactual Experiments: Homotopy method

- We only assume that there is such a mechanism, and that it is a smooth function of $\boldsymbol{\theta}$.
- Let $\boldsymbol{\pi}(\boldsymbol{\theta})$ be the (unique) selected equilibrium, for given $\boldsymbol{\theta}$, if we apply the "true" selection mechanism.
- Since we do not know the mechanism, we do not know $\boldsymbol{\pi}(\boldsymbol{\theta})$ for every possible $\boldsymbol{\theta}$.
- However, we DO know $\boldsymbol{\pi}(\boldsymbol{\theta})$ at the true $\boldsymbol{\theta}_0$ because we know that:

$$\mathbf{P}_0 = \boldsymbol{\pi}(\boldsymbol{\theta}_0)$$

and both $\mathbf{P}_0$ and $\boldsymbol{\theta}_0$ are identified.

Counterfactual Experiments: Homotopy method

- Let θ_0 and $\mathbf{P}_0$ be the the population values. Let $(\hat{\theta}_0, \hat{\mathbf{P}}_0)$ be our consistent estimator.
- We do not know the function $\pi(\theta)$. All what we know is that the point $(\hat{\theta}_0, \hat{\mathbf{P}}_0)$ belongs to the graph of this function π.
- Let θ^* be the vector of parameters under a counterfactual scenario.
- We want to know the counterfactual equilibrium $\pi(\theta^*)$.

Counterfactual Experiments: Homotopy method

- A Taylor approximation to $\pi(\theta^*)$ around our estimator $\hat{\theta}_0$ implies that:

$$\begin{aligned}\pi(\theta^*) &= \pi\left(\hat{\theta}_0\right) + \frac{\partial \pi\left(\hat{\theta}_0\right)}{\partial \theta'}\left(\theta^* - \hat{\theta}_0\right) + O\left(\left\|\theta^* - \hat{\theta}_0\right\|^2\right) \\ &= \hat{\mathbf{P}}_0 + \frac{\partial \pi\left(\hat{\theta}_0\right)}{\partial \theta'}\left(\theta^* - \hat{\theta}_0\right) + O\left(\left\|\theta^* - \hat{\theta}_0\right\|^2\right)\end{aligned}$$

- To get a first-order approximation to $\pi(\theta^*)$ we need to know $\dfrac{\partial \pi\left(\hat{\theta}_0\right)}{\partial \theta'}$.

Counterfactual Experiments: Homotopy method

- We know that $\pi\left(\hat{\boldsymbol{\theta}}_0\right) = \Psi(\hat{\boldsymbol{\theta}}_0, \hat{\mathbf{P}}_0)$, and this implies that:

$$\frac{\partial \pi\left(\hat{\boldsymbol{\theta}}_0\right)}{\partial \boldsymbol{\theta}'} = \left(I - \frac{\partial \Psi(\hat{\boldsymbol{\theta}}_0, \hat{\mathbf{P}}_0)}{\partial \mathbf{P}'}\right)^{-1} \frac{\partial \Psi(\hat{\boldsymbol{\theta}}_0, \hat{\mathbf{P}}_0)}{\partial \boldsymbol{\theta}'}$$

- Then, $\pi(\boldsymbol{\theta}^*) =$

$$\hat{\mathbf{P}}_0 + \left(I - \frac{\partial \Psi(\hat{\boldsymbol{\theta}}_0, \hat{\mathbf{P}}_0)}{\partial \mathbf{P}'}\right)^{-1} \frac{\partial \Psi(\hat{\boldsymbol{\theta}}_0, \hat{\mathbf{P}}_0)}{\partial \boldsymbol{\theta}'} \left(\boldsymbol{\theta}^* - \hat{\boldsymbol{\theta}}_0\right) + O\left(\left\|\boldsymbol{\theta}^* - \hat{\boldsymbol{\theta}}_0\right\|^2\right)$$

- Therefore, $\hat{\mathbf{P}}_0 + \left(I - \frac{\partial \Psi(\hat{\boldsymbol{\theta}}_0, \hat{\mathbf{P}}_0)}{\partial \mathbf{P}'}\right)^{-1} \frac{\partial \Psi(\hat{\boldsymbol{\theta}}_0, \hat{\mathbf{P}}_0)}{\partial \boldsymbol{\theta}'} \left(\boldsymbol{\theta}^* - \hat{\boldsymbol{\theta}}_0\right)$ is a first-order approximation to the counterfactual equilibrium $\mathbf{P}^*$.