INCAE PhD SUMMER ACADEMY
DYNAMIC GAMES IN EMPIRICAL IO

Lecture 1:
Single-agent dynamic discrete choice: Model and solution

Victor Aguirregabiria (University of Toronto)

June 20, 2022

ORGANIZATION OF THE COURSE

- This course is an Introduction to **dynamic discrete choice models in Empirical Industrial Organization**.

- It is organized in **5 Lectures** (from 9:00am to 10:20am) and **5 Tutorials** (from 11:00am to 12:20pm).

- **Lectures:** Models, methods, and some applications.

- **Tutorials:** Code for solution & estimation using real data.

- In the Tutorials, I will use GAUSS programming language, but you are encourage to use **your favorite programming language**, e.g., Matlab, R, Julia, Python, etc.

A BRIEF DESCRIPTION OF THE COURSE (1/2)

- This courses deals with **dynamic games of firms' competition** in Empirical IO.

- These models/applications deal with **firms' decisions** that:

  - Involve substantial **uncertainty**
  - Have effects on own firm's future profit (**dynamics**)
  - Have effects on competitors' profits (**game**)

- Some examples of applications are:
  - Market entry / exit
  - Investment in R&D, innovation
  - Investment in capacity, physical capital
  - Product design / quality
  - Pricing
  - Mergers
  - Networks (airlines, retail), ...

A BRIEF DESCRIPTION OF THE COURSE (2/2)

A. Before we study Dynamic Games, we will start with **Single-Agent Dynamic Discrete Choice models**:

- Model and solution (**Lecture 1**)
- Estimation methods (**Lecture 2**)

B. Then, we will extend this framework to allow for multiple players and strategic interactions, that is, **Dynamic Games**:

- Model and solution (**Lecture 3**)
- Estimation methods (**Lecture 4**)

C. We will conclude with **some applications**:

- Applications to firms' investment decisions (**Lecture 5**)

LECTURE 1: Single-Agent Dynamic Discrete Choice models

1. Introduction to Dynamic Structural Models
2. A Model of Market Entry / Exit
3. Optimal Decision Rules & Dynamic Programming
4. Conditional Choice Probabilities (CCPs)
5. Solution Methods

1. Introduction to Dynamic Discrete Choice Structural Models

SOME GENERAL FEATURES: DECISION & STATES

- t represents **time**, and it is discrete: $t \in \{1, 2, ...\}$.

- They are **econometric models** with a dependent variable y_t, explanatory variables $\mathbf{x}_t$, and unobservables to the researcher ε_t.

- $y_t =$ **agent's decision** at time t. It is discrete: $y_t \in \{0, 1, ..., J\}$

- The agent takes this action to **maximize her expected and discounted flow of utility** (payoffs):

$$\mathbb{E}_t \left(\sum_{s=0}^{T-t} \delta^s \ \pi_{t+s} \right)$$

$\delta \in [0, 1)$ is the discount factor, and π_t is the utility at period t.

SOME GENERAL FEATURES: UTILITY

- **Utility** depends on action y_t, and state variables $\mathbf{x}_t$ and ε_t:

$$\pi_t \;=\; \pi\left(y_t, \mathbf{x}_t, \varepsilon_t, \boldsymbol{\theta}_\pi\right)$$

  and $\boldsymbol{\theta}_\pi$ is the vector of **structural parameters** in the utility function.

- State variables in $\mathbf{x}_t$ are **observable to us as researchers**.

- State variables in ε_t are **unobservable to us as researchers**.

- Both $\mathbf{x}_t$ and ε_t are **known to the agent** at time t.

- Because the model is **dynamic**, $\mathbf{x}_t$ should depend somehow **on previous decisions**, y_{t-1} or/and y_{t-2}, ...

SOME GENERAL FEATURES: TRANSITIONS

- The model is completed with the specification of the **transition rules** or **transition probabilities** followed by the state variables $\mathbf{x}_t$ and ε_t.

- For ε_t, the most standard assumption is that it is i.i.d.

- For $\mathbf{x}_t$, the standard assumption is that it follows a **First Order Markov Process** that may depend on y_t with transition probability function:

$$Pr\left(\mathbf{x}_{t+1} = \mathbf{x}' \mid y_t = y, \ \mathbf{x}_t = \mathbf{x}\right) \ = \ f_x\left(\mathbf{x}' \mid y, \ \mathbf{x}; \boldsymbol{\theta}_f\right)$$

where $\boldsymbol{\theta}_f$ is a vector of structural parameters.

SOME GENERAL FEATURES: DYNAMIC PROGRAMMING

- The agent's decision problem is a **Dynamic Programming (DP)** problem.

- Let $V_t(\mathbf{x}_t, \varepsilon_t)$ be the value function at period t. The **Bellman Equation** of this DP problem is:

$$V_t(\mathbf{x}_t, \varepsilon_t) =$$

$$\max_{y_t} \left\{ \pi(y_t, \mathbf{x}_t, \varepsilon_t) + \delta \int V_{t+1}(\mathbf{x}_{t+1}, \varepsilon_{t+1}) f_x(d\mathbf{x}_{t+1} | y_t, \mathbf{x}_t) f_\varepsilon(d\varepsilon_{t+1}) \right\}$$

- The **Optimal Decision Rule** at period t, $\alpha_t(\mathbf{x}_t, \varepsilon_t)$, is the **argmax** in y_t of the expression within brackets $\{\}$.

SOME GENERAL FEATURES: MAIN PURPOSE

- Empirical applications of these models have two **main purposes**.

1. **Estimation of parameters in the utility function**.

   - Some of these parameters cannot be estimated from other sources and require a **Revealed Preference approach**

2. **Counterfactual experiments**.

   - We can use the estimated model to predict **agents' behavior in a counterfactual scenario** such as a new policy, or a change in some structural parameters.

2. Example:
Model of Market Entry & Exit

EXAMPLE: MARKET ENTRY & EXIT (1/2)

- Every period t, a firm decides whether to be active $(y_t = 1)$ or inactive $(y_t = 0)$ in a market.

- The **profit (utility) function** is:

$$\pi_t = \begin{cases} 0 & \text{if} \quad y_t = 0 \\ \\ \theta_1 + \theta_2 \; s_t + \theta_3 \; (1 - y_{t-1}) - \varepsilon_t & \text{if} \quad y_t = 1 \end{cases}$$

$s_t =$ market size, e.g., population, or average income in the market.

- **Economic interpretation**:
  $\theta_2 \; s_t =$ Variable profit.
  $-\theta_1 =$ Fixed cost.
  $-\theta_3 =$ Entry cost, i.e., extra if firm was inactive at previous period.
  $\varepsilon_t =$ mean zero shock in fixed costs.

EXAMPLE: MARKET ENTRY & EXIT (2/2)

- The vector of **observable state variable** $\mathbf{x}_t = (s_t, y_{t-1})$.

- s_t follows an AR(1) process and we represent its transition probability as $f_s(s_{t+1}|s_t)$.

- The **unobservable variable** ε_t is i.i.d. Logistic.

- The **Time horizon** T is infinite. $\delta \in [0, 1)$ is the discount factor.

- For the rest of this lecture, we use this simple model to present different concepts on dynamic discrete choice structural models.

3. Optimal Decision Rules & Dynamic Programming

BELLMAN EQUATION: STATIONARY MODEL

- When $T = \infty$, and functions π and f_x do not vary over time, the DP model is **stationary**, and value function $V(.)$ and optimal decision rule $\alpha(.)$ do not vary over time (**Blackwell Theorem**).

- Bellman equation can be written as:

$$V(\mathbf{x}_t, \varepsilon_t) = \max_{y_t \in \{0,1\}} \{y_t \ \bar{\pi}(\mathbf{x}_t) - y_t \ \varepsilon_t + \delta \ EV(y_t, s_t)\}$$

- $\bar{\pi}(\mathbf{x}_t)$ is the part of profit that does not depend on ε_t: i.e., $\bar{\pi}(\mathbf{x}_t) = \theta_1 + \theta_2 \ s_t + \theta_3 \ (1 - y_{t-1})$.

- $EV(y_t, s_t)$ is the **Continuation Value**:

$$EV(y_t, s_t) \ = \ \int V(y_t, s_{t+1}, \varepsilon_{t+1}) \ f_s(ds_{t+1}|s_t) \ f_\varepsilon(d\varepsilon_{t+1})$$

OPTIMAL DECISION RULE

- By definition of optimal decision rule, we have:

$$\alpha(\mathbf{x}_t, \varepsilon_t) = 1\{\bar{\pi}(\mathbf{x}_t) - \varepsilon_t + \delta\ EV(1, s_t) > \delta\ EV(0, s_t)\}$$

where $1\{.\}$ is the indicator function.

- Or equivalently:

$$\alpha(\mathbf{x}_t, \varepsilon_t) = 1\{\varepsilon_t < \bar{\pi}(\mathbf{x}_t) + \delta\ [EV(1, s_t) - EV(0, s_t)]\}$$

- The optimal decision rule **does not have a closed-form** analytical expression because function $EV(.)$ does not.

- i.e., need to solve numerically for the value function using Bellman eq.

INTEGRATED BELLMAN EQUATION (1/3)

- In this class of models, where ε_t is:
  (i) **not serially correlated**;
  (ii) **additive in the utility function**
  (iii) Logistic
  it is possible and computationally convenient to solve for the optimal decision rule using the **Integrated Bellman Equation**.

- We describe here the derivation of the Integrated Bellman Equation.

- First, define the **Integrated Value Function**:

$$V^{\sigma}(\mathbf{x}_t) \equiv \int V(\mathbf{x}_t, \varepsilon_t) \; f_{\varepsilon}(d\varepsilon_t)$$

INTEGRATED BELLMAN EQUATION (2/3)

- By definition, there is the following relationship between EV and V^{σ}:

$$EV(y_t, s_t) \;=\; \int V^{\sigma}(y_t, s_{t+1}) \; f_s(ds_{t+1}|s_t)$$

- Therefore, obtaining V^{σ} is sufficient to get $EV(.)$ and $\alpha(.)$.

- Integrating both sides of Bellman equation over the distribution of ε_t:

$$V^{\sigma}(\mathbf{x}_t) = \int \max_{y_t \in \{0,1\}} \{y_t \; \bar{\pi}(\mathbf{x}_t) - y_t \; \varepsilon_t + \delta \; EV(y_t, s_t)\} \;\; f_{\varepsilon}(d\varepsilon_t)$$

- This equation, together with the equation that relates EV and V^{σ} defines a **fixed point mapping** for V^{σ}. This mapping in the **Integrated Bellman Equation**

INTEGRATED BELLMAN EQUATION (3/3)

- When ε_t is Logistic (i.e., Type 1 Extreme Value), *Emax* operator has a closed form expression: i.e., the logarithm of the sum of exponentials.

- This implies the following form of the Integrated Bellman eq.:

$$V^{\sigma}(y_{t-1}, s_t) = \log\left(\exp\left\{\delta\ EV(0, s_t)\right\} + \exp\left\{\bar{\pi}(\mathbf{x}_t) + \delta\ EV(1, s_t)\right\}\right)$$

with

$$EV(y_t, s_t) \ = \ \int V^{\sigma}(y_t, s_{t+1})\ f_s(ds_{t+1}|s_t)$$

4. Conditional Choice Probabilities

CONDITIONAL CHOICE PROBABILITIES

- A key prediction of this model is the probability distribution of y_t conditional on $\mathbf{x}_t$.
- We denote this distribution as the **Conditional Choice Probability (CCP) function**.
- More precisely, for any value of $(y, \mathbf{x})$, the CCP function $P(y|\mathbf{x})$ is defined as:

$$P(y|\mathbf{x}) \equiv \Pr\left(\alpha(\mathbf{x}_t, \varepsilon_t) = y \mid \mathbf{x}_t = \mathbf{x}\right)$$

- These are the model predictions that we use to estimate the parameters of the model.

CONDITIONAL CHOICE PROBABILITIES (2/2)

- Remember that:

$$\alpha(\mathbf{x}_t, \varepsilon_t) \;=\; 1\left\{\varepsilon_t \;<\; \bar{\pi}(\mathbf{x}_t) + \delta\; \left[EV(1, s_t) - EV(0, s_t)\right]\right\}$$

- Therefore, if ε_t is Logistic, we have that:

$$P(1|\mathbf{x}_t) = \frac{\exp\left\{\bar{\pi}(\mathbf{x}_t) + \delta\; \left[EV(1, s_t) - EV(0, s_t)\right]\right\}}{1 + \exp\left\{\bar{\pi}(\mathbf{x}_t) + \delta\; \left[EV(1, s_t) - EV(0, s_t)\right]\right\}}$$

- These are the probabilities we use in the estimation of the model by Maximum Likelihood or other methods.

5. Solution Methods

FIXED-POINT ITERATIONS INTEGRATED BELLMAN EQ.

- I describe a Fixed-Point algorithm in the Integrated Bellman equation to obtain the V^{σ} and the solution of the model.

- I describe it in **Vector Form** for easy implementation in **Vector Programming Languages** such as Gauss, Matlab, R, Julia, Python.

- Suppose that s_t is discrete: $s_t \in \{s^1, s^2, ..., s^{|S|}\}$.

- The primitives of the model are:

1. **Vectors of payoffs**: $\mathbf{\Pi}(0)$ and $\mathbf{\Pi}(1)$ with dimension $|S| \times 1$, with $\bar{\pi}(0, s)$ and $\bar{\pi}(1, s)$ for every value of s.

2. **Matrix of transition probabilities**: $\mathbf{F}_s$ with dimension $|S| \times |S|$, with probabilities $f_s(s_{t+1}|s_t)$.

3. **Discount factor**: δ.

FIXED-POINT ITERATIONS IN INT. BELLMAN EQ. (2/3)

- We can represent the value function $V^{\sigma}(y_{t-1}, s_t)$ in terms of two vectors, $\boldsymbol{V}^{\sigma}(0)$ and $\boldsymbol{V}^{\sigma}(1)$, each with dimension $|S| \times 1$.

- Given $\boldsymbol{V}^{\sigma}(0)$ and $\boldsymbol{V}^{\sigma}(1)$ and transition matrix $\mathbf{F}_s$, we can represent the continuation value function $EV(y_t, s_t)$ using two vectors:

$$\mathbf{F}_s \ \boldsymbol{V}^{\sigma}(0) \qquad and \qquad \mathbf{F}_s \ \boldsymbol{V}^{\sigma}(1)$$

- Then, we can represent the Integrated Bellman equation in terms of two vector-valued equations:

$$\begin{cases} \boldsymbol{V}^{\sigma}(0) & = & \log\left(\exp\left\{\delta \ \mathbf{F}_s \ \boldsymbol{V}^{\sigma}(0)\right\} + \exp\left\{\mathbf{\Pi}(0) + \delta \ \mathbf{F}_s \ \boldsymbol{V}^{\sigma}(1)\right\}\right) \\ \boldsymbol{V}^{\sigma}(1) & = & \log\left(\exp\left\{\delta \ \mathbf{F}_s \ \boldsymbol{V}^{\sigma}(0)\right\} + \exp\left\{\mathbf{\Pi}(1) + \delta \ \mathbf{F}_s \ \boldsymbol{V}^{\sigma}(1)\right\}\right) \end{cases}$$

FIXED-POINT ITERATIONS (3/3)

- We start with (arbitrary) initial vectors $\boldsymbol{V}_0^{\sigma}(0)$ and $\boldsymbol{V}_0^{\sigma}(1)$.

- At each iteration $n \geq 1$ we update these vectors and obtain $\boldsymbol{V}_n^{\sigma}(0)$ and $\boldsymbol{V}_n^{\sigma}(1)$ using:

$$\begin{cases} \boldsymbol{V}_n^{\sigma}(0) & = & \log\left(\exp\left\{\delta\ \mathbf{F}_s\ \boldsymbol{V}_{n-1}^{\sigma}(0)\right\} + \exp\left\{\boldsymbol{\Pi}(0) + \delta\ \mathbf{F}_s\ \boldsymbol{V}_{n-1}^{\sigma}(1)\right\}\right. \\ \boldsymbol{V}_n^{\sigma}(1) & = & \log\left(\exp\left\{\delta\ \mathbf{F}_s\ \boldsymbol{V}_{n-1}^{\sigma}(0)\right\} + \exp\left\{\boldsymbol{\Pi}(1) + \delta\ \mathbf{F}_s\ \boldsymbol{V}_{n-1}^{\sigma}(1)\right\}\right. \end{cases}$$

- We reach convergence at iteration n if $||\boldsymbol{V}_n^{\sigma}(0) - \boldsymbol{V}_{n-1}^{\sigma}(0)||$ and $||\boldsymbol{V}_n^{\sigma}(1) - \boldsymbol{V}_{n-1}^{\sigma}(1)||$ are both smaller that a pre-specified small constant close to zero (e.g., 1e-6).

- The integrated Bellman equation is a contraction mapping. This implies that Fixed Point iterations converge to the unique solution.