Dynamic Games: Numerical Methods and Applications

Mini Course, Tel Aviv University, Day 1
Prof. Ulrich Doraszelski, University of Pennsylvania

What are Dynamic Games?

- A tool for analyzing dynamic strategic interactions.
  - dynamic $\rightarrow$ forward-looking players optimize over time;
  - strategic $\rightarrow$ each player recognizes that its actions impact other players.

- Often used to track evolution of oligopolistic industries.
  - oligopolistic $\rightarrow$ neither perfectly competitive nor monopolistically competitive.

Dynamics + Strategic Interactions = Dynamic Games

- Combine literature on long-run industry equilibrium (Jovanovic 1982, Hopenhayn 1992, Melitz 2003) with game theory (Tirole 1988, Fudenberg & Tirole 1991).

Why use Dynamic Games?

- Key findings of empirical literature on industry evolution (Mueller 1986, Dunne, Roberts, & Samuelson 1988, Davis & Haltiwanger 1992):
  - Entry and exit occur simultaneously.
  - Heterogeneity among firms evolves endogenously in response to random occurrences.
  - Heterogeneity among firms persists over long stretches of time.

Why use Dynamic Games?

- Game theory revolution in economics: emphasis on analytically tractable models.
  - End effects.
  - Transitional dynamics.
  - Inherently dynamic phenomena.

Agenda

- From dynamic programming to dynamic games.

- Application: Quality ladder model without entry/exit.

From Dynamic Programming...

- Time is discrete. The horizon is infinite.

- The state space $\Omega = \{1, 2, \ldots, L\}$ is finite.

- The state in period t is $\omega_t \in \Omega$. The law of motion is a controlled discrete-time, finite-state, first-order Markov process, where
$$\Pr(\omega_{t+1}|\omega_t, x_t)$$
is the probability that the state transits from ω_t to ω_{t+1} if the control is $x_t \in D(\omega_t)$ and $D(\omega_t)$ is the nonempty set of feasible controls in state ω_t.

- The objective is to maximize the expected NPV of payoffs
$$E\left\{\sum_{t=0}^{\infty} \beta^t \pi(\omega_t, x_t)\right\},$$
where $\beta \in [0, 1)$ is the discount factor and $\pi(\omega_t, x_t)$ is the per-period payoff in state ω_t if the control is x_t.

- The value function $V(\omega)$ is the maximum expected NPV of present and future payoffs if the current state is ω. It satisfies the Bellman equation
$$V(\omega) = \max_{x \in D(\omega)} \pi(\omega, x) + \beta \sum_{\omega'=1}^{L} V(\omega')\Pr(\omega'|\omega, x) \tag{1}$$
and the optimal policy function $X(\omega)$ satisfies
$$X(\omega) \in \arg\max_{x \in D(\omega)} \pi(\omega, x) + \beta \sum_{\omega'=1}^{L} V(\omega')\Pr(\omega'|\omega, x).$$

- The collection of equation (1) for all states $\omega \in \Omega$ defines a system of nonlinear equations. The contraction mapping theorem ensures existence and uniqueness of a solution.

... to Dynamic Games

- N players.

- The law of motion is a controlled discrete-time, finite-state, first-order Markov process, where
$$\text{Pr}(\omega_{t+1}|\omega_t, x_t)$$
is the probability that the state transits from ω_t to ω_{t+1} if the controls are $x_t = (x_{1t}, \ldots, x_{Nt}) \in \times_{n=1}^{N} D_n(\omega_t)$ and $D_n(\omega_t)$ is the nonempty set of feasible controls of player n in state ω_t.

- $\pi_n(\omega_t, x_t)$ is the per-period payoff of player n in state ω_t if the controls are x_t.

- The value function $V_n(\omega)$ of player n satisfies the Bellman equation
$$V_n(\omega) = \max_{x_n \in D_n(\omega)} \pi_n(\omega, x_n, X_{-n}(\omega)) + \beta \sum_{\omega'=1}^{L} V_n(\omega')\text{Pr}(\omega'|\omega, x_n, X_{-n}(\omega)) \tag{2}$$
and his optimal policy function $X_n(\omega)$ satisfies
$$X_n(\omega) \in \arg\max_{x_n \in D_n(\omega)} \pi_n(\omega, x_n, X_{-n}(\omega)) + \beta \sum_{\omega'=1}^{L} V_n(\omega')\text{Pr}(\omega'|\omega, x_n, X_{-n}(\omega)). \tag{3}$$

- The collection of equations (2) and (3) for all states $\omega \in \Omega$ and all players $n = 1, \ldots, N$ defines a Markov-perfect equilibrium. The contraction mapping theorem does not apply and neither existence nor uniqueness of a MPE is guaranteed.

. . . to Dynamic Games

- Special case: ω is a vector partitioned into
$$(\omega_1, \ldots, \omega_N),$$
where ω_n denotes the (one or more) coordinates of the state that describe player n.
Examples: Production capacity, marginal cost, product quality.
Nomenclature:
  - $\omega_n \in \Omega_n = \{1, 2, \ldots, L_n\}$ is the state of player n;
  - $\omega \in \times_{n=1}^{N} \Omega_n$ is the state of the game.

  Equations (2) and (3) can be written as
$$V_n(\omega) = \max_{x_n \in D_n(\omega)} \pi_n(\omega, x_n, X_{-n}(\omega)) + \beta \sum_{\omega_1'=1}^{L_1} \ldots \sum_{\omega_N'=1}^{L_N} V_n(\omega') \mathrm{Pr}(\omega' | \omega, x_n, X_{-n}(\omega)),$$
$$X_n(\omega) \in \arg\max_{x_n \in D_n(\omega)} \pi_n(\omega, x_n, X_{-n}(\omega)) + \beta \sum_{\omega_1'=1}^{L_1} \ldots \sum_{\omega_N'=1}^{L_N} V_n(\omega') \mathrm{Pr}(\omega' | \omega, x_n, X_{-n}(\omega)).$$

- Even more special case: Transitions in player n's state are controlled by player n's actions and are independent of the actions of other players and transitions in their states, i.e.,
$$\mathrm{Pr}\left(\omega' | \omega, x\right) = \prod_{n=1}^{N} \mathrm{Pr}_n\left(\omega_n' | \omega_n, x_n\right).$$

Quality Ladder Model without Entry/Exit

- Pakes, A. & McGuire, P. (1994) "Computing Markov-Perfect Nash Equilibria: Numerical Implications of a Dynamic Differentiated Product Model."
- Borkovsky, R., Doraszelski, U. & Kryukov, Y. (2010) "A User's Guide to Solving Dynamic Stochastic Games Using the Homotopy Continuation Method."
- Discrete time, infinite horizon.
- Two firms with potentially different product qualities

$$\omega = (\omega_1, \omega_2) \in \{1, \ldots, L\}^2 = \Omega.$$

- In each period, the timing is as follows:
  - Firms choose investments in quality improvements.
  - Product market competition takes place.
  - Investment outcomes and depreciation shocks are realized.

Product Market Competition

- Firm n's demand is

$$D_n(p_1, p_2; \omega) = M \frac{\exp\left(g(\omega_n) - p_n\right)}{1 + \sum_{k=1}^{2} \exp\left(g(\omega_k) - p_k\right)},$$

  where $M > 0$ is market size and

$$g(\omega_n) = \begin{cases} 3\omega_n - 4 & \text{if} \quad \omega_n \leq 5, \\ 12 + \ln\left(2 - \exp\left(16 - 3\omega_n\right)\right) & \text{if} \quad \omega_n > 5 \end{cases}$$

  maps product quality into consumers' valuations.

- Firm n solves

$$\max_{p_n \geq 0} D_n(p_1, p_2; \omega)(p_n - c),$$

  where c is marginal cost of production.

- FOC:

$$0 = 1 - \frac{1 + \exp\left(g(\omega_{-n}) - p_{-n}\right)}{1 + \exp\left(g(\omega_n) - p_n\right) + \exp\left(g(\omega_{-n}) - p_{-n}\right)}(p_n - c), \quad n \neq -n.$$

- Compute Nash equilibrium $(p_1(\omega), p_2(\omega))$ by solving system of FOCs.

- Firm n's profit is

$$\pi_n(\omega) = D_n(p_1(\omega), p_2(\omega); \omega)(p_n(\omega) - c).$$

Investment Dynamics

- Let $x_n \geq 0$ be firm n's investment in quality improvements.
- Law of motion:
  - Successful investment has probability $\frac{\alpha x_n}{1+\alpha x_n}$.
  - Depreciation shock has probability δ.
- Transition probability: If $\omega_n \in \{2, \ldots, L-1\}$, then

$$\Pr(\omega_n'|\omega_n, x_n) = \begin{cases} \frac{(1-\delta)\alpha x_n}{1+\alpha x_n} & \text{if} \quad \omega_n' = \omega_n + 1, \\ \frac{1-\delta+\delta\alpha x_n}{1+\alpha x_n} & \text{if} \quad \omega_n' = \omega_n, \\ \frac{\delta}{1+\alpha x_n} & \text{if} \quad \omega_n' = \omega_n - 1. \end{cases}$$

  If $\omega_n \in \{1, L\}$, then

$$\Pr(\omega_n'|1, x_n) = \begin{cases} \frac{(1-\delta)\alpha x_n}{1+\alpha x_n} & \text{if} \quad \omega_n' = 2, \\ \frac{1+\delta\alpha x_n}{1+\alpha x_n} & \text{if} \quad \omega_n' = 1, \end{cases}$$

$$\Pr(\omega_n'|L, x_n) = \begin{cases} \frac{1-\delta+\alpha x_n}{1+\alpha x_n} & \text{if} \quad \omega_n' = L, \\ \frac{\delta}{1+\alpha x_n} & \text{if} \quad \omega_n' = L - 1. \end{cases}$$

Bellman Equation

- Let $V_n(\omega)$ denote the expected NPV to firm n if the current state is ω.

- Firm n's Bellman equation is

$$V_n(\omega) = \max_{x_n \geq 0} \pi_n(\omega) - x_n + \beta \sum_{\omega'_n=1}^{L} W_n(\omega'_n; \omega_{-n}, x_{-n}(\omega)) \mathrm{Pr}(\omega'_n | \omega_n, x_n),$$

  where

  - the expectation (with respect to its rival's successor state) of firm n's continuation value in state ω'_n is

$$W_n(\omega'_n; \omega_{-n}, x_{-n}(\omega)) = \sum_{\omega'_{-n}=1}^{L} V_n(\omega') \mathrm{Pr}(\omega'_{-n} | \omega_{-n}, x_{-n}(\omega));$$

  - $x_{-n}(\omega)$ is the rival's investment strategy;

  - $\beta \in [0, 1)$ is the discount factor.

Investment Strategy

- Firm n's investment strategy is

$$x_n(\omega) = \arg\max_{x_n \geq 0} \pi_n(\omega) - x_n + \beta \sum_{\omega_n'=1}^{L} W_n(\omega_n')\Pr(\omega_n'|\omega_n, x_n),$$

where $W_n(\omega_n')$ is shorthand for $W_n(\omega_n'; \omega_{-n}, x_{-n}(\omega))$.

- If $\omega_n \in \{2, \ldots, L-1\}$, then

$$x_n(\omega) = \frac{-1 + \sqrt{\max\{1, \beta\alpha\left((1-\delta)(W_n(\omega_n+1) - W_n(\omega_n)) + \delta(W_n(\omega_n) - W_n(\omega_n-1))\right)\}}}{\alpha}.$$

If $\omega_n \in \{1, L\}$, then

$$x_n(\omega) = \frac{-1 + \sqrt{\max\{1, \beta\alpha(1-\delta)\left(W_n(2) - W_n(1)\right)\}}}{\alpha},$$

$$x_n(\omega) = \frac{-1 + \sqrt{\max\{1, \beta\alpha\delta\left(W_n(L) - W_n(L-1)\right)\}}}{\alpha}.$$

Equilibrium

- Profits from product market competition are symmetric:

$$\pi_1(\omega_1, \omega_2) = \pi_2(\omega_2, \omega_1).$$

  The remaining primitives are also symmetric.

- Symmetric Markov perfect equilibrium (MPE):
  - Value function $V_1(\omega_1, \omega_2) = V(\omega_1, \omega_2)$ and $V_2(\omega_1, \omega_2) = V(\omega_2, \omega_1)$.
  - Policy function $x_1(\omega_1, \omega_2) = x(\omega_1, \omega_2)$ and $x_2(\omega_1, \omega_2) = x(\omega_2, \omega_1)$.

- Existence in pure strategies is guaranteed (Doraszelski & Satterthwaite 2010), uniqueness is not.

- The goal is to compute the value and policy functions (or, more precisely, $L \times L$ matrices) $\mathbf{V}$ and $\mathbf{x}$.

Computation: Pakes & McGuire (1994) Algorithm

1. Make initial guesses $\mathbf{V}^0$ and $\mathbf{x}^0$, choose a stopping criterion $\epsilon > 0$, and initialize the iteration counter to $k = 1$.

2. For all states $\omega \in \Omega$ compute

$$x^{k+1}(\omega) = \arg\max_{x_1 \geq 0} \pi_1(\omega) - x_1 + \beta \sum_{\omega_1'=1}^{L} W^k(\omega_1') \Pr(\omega_1' | \omega_1, x_1)$$

   and

$$V^{k+1}(\omega) = \pi_1(\omega) - x^{k+1}(\omega) + \beta \sum_{\omega_1'=1}^{L} W^k(\omega_1') \Pr(\omega_1' | \omega_1, x^{k+1}(\omega)),$$

   where

$$W^k(\omega_1') = \sum_{\omega_2'=1}^{L} V^k(\omega') \Pr(\omega_2' | \omega_2, x^k(\omega_2, \omega_1)).$$

3. If

$$\max_{\omega \in \Omega} \left| \frac{V^{k+1}(\omega) - V^k(\omega)}{1 + |V^{k+1}(\omega)|} \right| < \epsilon \quad \wedge \quad \max_{\omega \in \Omega} \left| \frac{x^{k+1}(\omega) - x^k(\omega)}{1 + |x^{k+1}(\omega)|} \right| < \epsilon$$

   then stop; else increment the iteration counter k by one and go to step 2.