Lecture 5: Strategic Models of Industry Dynamics

Jean-François Houde
UW-Madison

December 12, 2023

Dynamic Investment Game with Entry/Exit

Reference: Ericson and Pakes (1995)

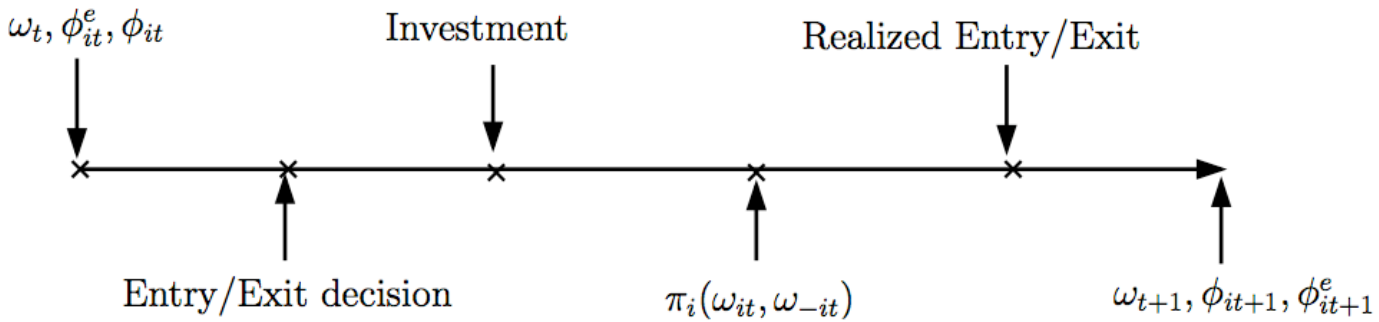
- **Introduction:**
  - Central element: Productivity is stochastically controlled by firms' investments.
  - Investment is strategic: Markov Perfect Equilibrium (MPE).
  - Model is general enough that it can be applied in various contexts:
    1. Quality ladder model (Pakes and McGuire (1994)),
    2. Learning-by-doing (Benkard (2004), Besanko et al. (2010)),
    3. Capacity dynamics (Besanko and Doraszelski (2004)),
    4. Horizontal mergers (Gowrisankaran (1999)),
    5. Advertisement (Doraszelski and Markovich (2005)).
- See Doraszelski and Pakes (2007) for a review.

Model

- **Actions:**
  1. Static actions: Price/quantity game. Lead to a (unique) static payoff function $\pi_i(\omega_i, \omega_{-i})$.
  2. Dynamic actions: Entry/Exit, Investment $x_i(\omega_i, \omega_{-i})$.
- **Assumptions:** $\pi_i(\omega_i, \omega_{-i})$ is continuous, bounded, and increasing in ω_i.
- **State space:**
  - Firm i's productivity: $\omega_i \in \Omega = \{1, 2, ..., \bar{\omega}\}$ (public information).
    $*$ ω_i can represents a quality index, a marginal cost shifter, capacity, etc.
  - Entry cost: $\phi_i^e \sim F^e(.)$ (private information).
  - Scrap value: $\phi_i \sim F(.)$ (private information).
  - Industry state: $\omega \in S = \{\omega_1, ..., \omega_n | \omega_i \in \Omega, n \leq \bar{n}\}$
  - **Note:** $|S|$ grows exponentially in $\bar{n}$ (curse of dimensionality).

Model (continued)

- **Assumptions:** Payoffs and strategies are **symmetric** and **anonymous**.
  - Symmetry: $f_i(\omega_i, \omega_{-i}) = f_j(\omega_i, \omega_{-i})$.
  - Anonymous: $f(\omega_i, \omega_{-i}) = f(\omega_i, \omega_{\mu(-i)})$.
  - Where $\mu(-i)$ is any reordering of indexes $-i$.
- The symmetry and anonymity assumptions reduce dramatically the industry state space. Two ways of writing the industry state:
  1. $\omega \in S^0 = \{\omega_1, ..., \omega_n | \omega_i \in \Omega, \omega_1 > \omega_2 > ... > \omega_n, n \leq \bar{n}\} \subset S$
  2. $\omega \in S^0 = \{s_1, ..., s_{\bar{\omega}} | \omega_i \in \Omega, \sum_\omega s_\omega \leq \bar{n}\} \subset S$
- **Timing:**



Strategic actions

- **Investment:**
  - Firms invest x_i to raise the probability of getting a productivity gain ν next period:
$$\Pr(\omega_i' = \omega_i + \nu_i | \omega_i, x_i) = P_\nu(\nu_i | x, \omega_i)$$
where $\nu_i \in \{0, 1\}$ and $p_\nu(\nu_i = 1 | x_i = 0, \omega_i) = 0$.
  - Assumption: $p_\nu(\nu = 1 | x, \omega)$ is strictly increasing in x_i.
  - Example (Pakes and McGuire (1994)):
$$p_\nu(\nu = 1 | x, \omega) = p_\nu(\nu = 1 | x) = \frac{\alpha x}{1 + \alpha x}$$
- **Exit:**
  - At the beginning of the period incumbents observe scrap value of equipments ϕ_i and decide to exit at the end of the period (i.e. all incumbents receive $\pi(\omega_i, \omega_{-i})$).
  - Since ϕ_i is private information, we can characterize the exit strategy either by an exit cutoff $\bar{\phi}(\omega_i, \omega_{-i})$, or by the continuation probability $r(\omega_i, \omega_{-i}) = \Pr(\phi < \bar{\phi}(\omega_i, \omega_{-i}))$.

Strategic actions (continued)

- **Entry:**
  - Each period there are $\mathcal{E}$ potential entrants who live **only** one period.
  - Potential entrants privately observe an entry cost ϕ^e and decides to enter or not at the end of the period (one period set-up cost).
  - Entrants choose to invest x^e in order to improve their productivity next period.
  - Entry strategy: Optimal cutoff point $\bar{\phi}^e(\omega_i, \omega_{-i})$, or entry probability $r^e = \Pr(\phi < \bar{\phi}^e(\omega_i, \omega_{-i}))$.
- **State-to-state transitions:**
  - Individual state transition:

$$\omega_i' = \omega_i + \underbrace{\nu_i}_{\text{productivity gain}} - \underbrace{\eta}_{\text{aggregate shock}},$$

    where $\eta \in \{0, 1\}$ and $\Pr(\eta = 1) = \delta$.
    Note: The depreciation shock η can either be industry-wide or firm specific.
  - Industry state transition probability matrix: $q(\omega'_{-i}|\omega, \eta)$ is the perceived transition probabilities of competitors' states (i.e. beliefs).

Value function: Incumbent

- Incumbent's problem prior to observing ϕ:

$$v(\omega_i, \omega_{-i}) = \pi(\omega_i, \omega_{-i}) + (1 - r_i)E\left[\phi | \beta \sum_{\nu} W_i(\nu|\omega) p_\nu(\nu|x_i) - x_i < \phi\right]$$
$$+ r_i\Big[- x_i + \beta \sum_{\nu} W_i(\nu|\omega) p_\nu(\nu|x_i)\Big]$$

- Where:
  - $W_i(\nu|\omega) = \sum_{\omega'_{-i},\eta} v(\omega_i + \nu - \eta, \omega'_{-i}) q(\omega'_{-i}|\omega_i, \omega_{-i}, \eta) p_\eta(\eta)$,
  - $\omega = (\omega_i, \omega_{-i})$.
  - r_i is the optimal continuation probability:

$$r_i(\omega, \omega_{-i}) = F\big(\beta \sum_{\nu} W_i(\nu|\omega) p_\nu(\nu|x_i) - x_i\big)$$

  - x_i solves the following KT-FOC:

$$x_i\left(\beta \sum_{\nu} W(\nu|\omega) \frac{\partial p_\nu(\nu|x_i)}{\partial x_i} - 1\right) = 0$$

Value function: Potential Entrant

- Entrant's problem:

$$V^e(\omega, \phi^e) = \max_{x^e} \left\{ 0, -\phi^e - x_i^e + \beta \sum_{\nu} W(\nu|\omega) p_\nu(\nu|x^e) \right\}$$

Where W is the continuation value of an entrant (i.e. starting at $\omega_i = \omega_e$), x^e is defined as before. The entry probability is given by:

$$r^e(\omega_i, \omega_{-i}) = F^e\Big(- x^e(\omega) + \beta \sum_{\nu} W(\nu|\omega) p_\nu(\nu|x^e(\omega)) \Big)$$

- **Equilibrium:** A Markov Perfect Equilibrium (MPE) equilibrium is a set of functions $\{x, r, x^e, r^e\}$ such that:
  1. The policy functions solve the incumbents and entrants problems given beliefs $q(.|.)$.
  2. The perceived aggregate transition probabilities $q(\omega'_{-i}|\omega, \eta)$ are consistent with the optimal response of all agents.

Properties of the equilibrium

- Existence of a pure strategy equilibrium:
  - Continuity of the scrap value and entry cost distributions ensure the existence of unique entry/exit strategies in probability space (i.e. perfect-bayesian equilibrium, see Doraszelski and Satterthwaite (2005) for the exact conditions of profits and F).
  - As long the $p_\nu(.|x)$ satisfies some regularity conditions such that the continuation value is concave, there exits a unique best-reply investment function (see Ericson and Pakes (1995)).
- Uniqueness: Very difficult to show, apart for some special cases.

Properties of the equilibrium

- Characterization:
  - There exists a pair $(\bar{n}, \bar{\omega})$ such that with probability one:
    - ⋆ there will never be more than $\bar{n}$ firms active,
    - ⋆ we will never observe an active firm with $\omega_i > \bar{\omega}$ (i.e. $x_i(\bar{\omega}, \omega_{-i}) = 0, \forall \omega_{-i}$).

    Therefore, the state space is always **finite** and we can compute an equilibrium.
- The equilibrium defines a finite state Markov chain with transition probability matrix $Q(\omega', \omega)$. Therefore there exists at least one recurrent class of states $R \in S^0$, such that:
  - with probability one each sample path $\{\omega_t\}$ will enter R,
  - and once in R the process will stay within it forever.

Note: Since R is typically much smaller and S^0, if we can identify it, we can reduce considerably the computation burden (see Pakes and McGuire (2001)).

Computing MPE Nash equilibrium: Gauss-Jacobi VS Gauss-Seidel algorithms

- **Example:**
  - Static entry game with 2 players and incomplete information.
  - Payoff of entering:

$$\pi_i = \alpha_i - \theta E(N_{-i}) + \epsilon_i = \alpha_i - \theta \tilde{p}_{-i} + \epsilon_i$$

  - If ϵ_i is iid extreme value, the best-response function is given by:

$$p_i = R(\tilde{p}_{-i}) = \frac{\exp(\alpha_i - \theta \tilde{p}_{-i})}{1 + \exp(\alpha_i - \theta \tilde{p}_{-i})}$$

    and a Nash equilibrium is such that $p_i = R(p_{-i})$ for all is.

- **Gauss-Jacobi**:

  it 0: Guess p_i^0.

  it k: Let $\tilde{p}_j = p_j^{k-1}$, and update strategies using the best-response mapping:

$$p_i^k = R(\tilde{p}_{-i})$$

  test: Stop if $||p^k - p^{k-1}|| < \delta$, else update $k + 1$.

Computing MPE Nash equilibrium: Gauss-Jacobi VS Gauss-Seidel algorithms

- **Gauss-Seidel**:

  it 0: Guess p_i^0.

  it $k.1$: Let $\tilde{p}_{-1} = p_{-1}^{k-1}$, and update strategies using the best-response mapping for player 1:

  $$p_1^k = R(\tilde{p}_{-1})$$

  ...

  it $k.j$: Let $\tilde{p}_i = p_i^{k-1}$ for $i > j$, and $\tilde{p}_i = p_i^k$ for $i < j$. Update strategies using the best-response mapping for player i:

  $$p_j^k = R(\tilde{p}_{-j})$$

  ...

  test: Stop if $||p^k - p^{k-1}|| < \delta$, else update $k + 1$.

- The Guass-Seidel usually converges faster, but both algorithm can converge to different equilibria (or not converge at all...).
- Alternatives:
  - Damping: $p_i^k = \alpha R(\tilde{p}_{-i}) + (1 - \alpha) p_i^{k-1}$
  - Random updating orders (in the Gauss-Seidel algorithm).

Besanko and Doraszelski (2004): Capacity dynamics and endogenous asymmetries in firm size

- **Stylized facts**
  - Market structure is persistent
  - Most industries exhibit important asymmetries: Dominant firm(s) competing with fringes.
  - How heterogeneous initial conditions emerge in the first place?
  - Cournot vs Bertrand: Equivalent market-structure?

Model

- Duopoly with homogenous products:

$$D(P) = a - bP \Rightarrow P(Q) = \frac{a}{b} - \frac{1}{b}(q_1 + q_2) = A - BQ$$

- Heterogeneous (discrete) capacities:

$$0 \leq q_i \leq \bar{q}_i$$

  and industry structure: $\omega_t = \{\bar{q}_{1t}, \bar{q}_{2t}\}$.
- Two models of competition:
  1. Quantity: Cournot
  2. Price competition: Edgeworth-Bertrand (see for instance Deneckere & Kovenock)
- Investment affects the prob. of capacity improvement

$$\Pr(\Delta\bar{q}_{1t}|\bar{q}_{1t}, x_{1t}) = \begin{cases} \frac{(1-\delta)\alpha x_{1t}}{1+\alpha x_{1t}} & \text{If } \Delta q_{1t} = 1, \\ \frac{(1-\delta)}{1+\alpha x_{1t}} + \frac{\delta\alpha x_{1t}}{1+\alpha x_{1t}} & \text{If } \Delta q_{1t} = 0, \\ \frac{\delta}{1+\alpha x_{1t}} & \text{If } \Delta q_{1t} = -1 \end{cases}$$

- This is in the interior of the capacity grid. Boundary states: $\bar{q} = M$ and $\bar{q} = 0$.

Cournot Duopoly Example

- Heterogeneous (discrete) capacities:

$$0 \leq q_i \leq \bar{q}_i$$

  and industry structure: $\omega_t = \{\bar{q}_{1t}, \bar{q}_{2t}\}$.
- Best-response functions (zero marginal cost):

$$\max_{q_i < \bar{q}_i} \quad (A - Bq_i - Bq_{-i})\, q_i$$

$$\rightarrow A - 2Bq_i - Bq_{-i} = 0 \quad \text{(interior solution)}$$

$$R(q_{-i}) = \min \left\{ \bar{q}_i, \frac{A}{2B} - \frac{q_{-i}}{2} \right\}$$

- Nash equilibrium: $\{q_1(\bar{q}_1, \bar{q}_2), q_2(\bar{q}_1, \bar{q}_2)\}$
  1. *Unconstrained:* $q_i^* = q_j^* = \frac{A}{3B}$ if $\bar{q}_i \leq \frac{A}{3B}$ for all i.
  2. *Partial constrained:* $q_i^* = R(\bar{q}_{-i})$ and $q_{-i}^* = \bar{q}_{-i}$ if $R(\bar{q}_{-i}) \leq \bar{q}_i$ and $R(q_i^*) > \bar{q}_{-i}$
  3. *Constrained:* $q_i^* = \bar{q}_i$ if $R(\bar{q}_{-i}) > \bar{q}_i$ and $R(\bar{q}_i) > \bar{q}_{-i}$
- Nash-equilibrium profit function:

$$\pi(\omega_t) = (A - Bq_1(\omega_t) - Bq_2(\omega_t)) \times q_1(\omega_t)$$

Investment decision

- Fix rival strategy: $\tilde{x}(\omega)$
- Bellman equation:

$$V_1^{\tilde{x}}(\omega) = \pi_1(\omega) + \max_{x_1 \geq 0} \left\{ -x_1 + \beta \sum_{\bar{q}_1'} W_1^{\tilde{x}}(\bar{q}_1'|\omega) \Pr(\bar{q}_1'|\bar{q}_1, x_1) \right\}$$

  where $W_1^{\tilde{x}}(\bar{q}_1) = \sum_{\bar{q}_2'} V_1^{\tilde{x}}(\bar{q}_1, \bar{q}_2) \Pr(\bar{q}_2'|\bar{q}_2, \tilde{x}_2)$.
- Kuhn-Tucker condition:

$$-1 + \beta \sum_{\bar{q}_1'} W_1^{\tilde{x}}(\bar{q}_1'|\omega) \frac{\partial \Pr(\bar{q}_1'|\bar{q}_1, x_1)}{\partial x_1} \geq 0$$

- Closed-form best-response function:

$$x_1(\omega) = R_1(\tilde{x}|\omega) = \max \left\{ 0, \frac{-1 + \sqrt{\beta\alpha \left((1-\delta) \left(W_1^{\tilde{x}}(\bar{q}_1 + 1) - W_1^{\tilde{x}}(\bar{q}_1) \right) + \delta \left(W_1^{\tilde{x}}(\bar{q}_1) - W_1^{\tilde{x}}(\bar{q}_1 - 1) \right) \right)}}{\alpha} \right\}$$

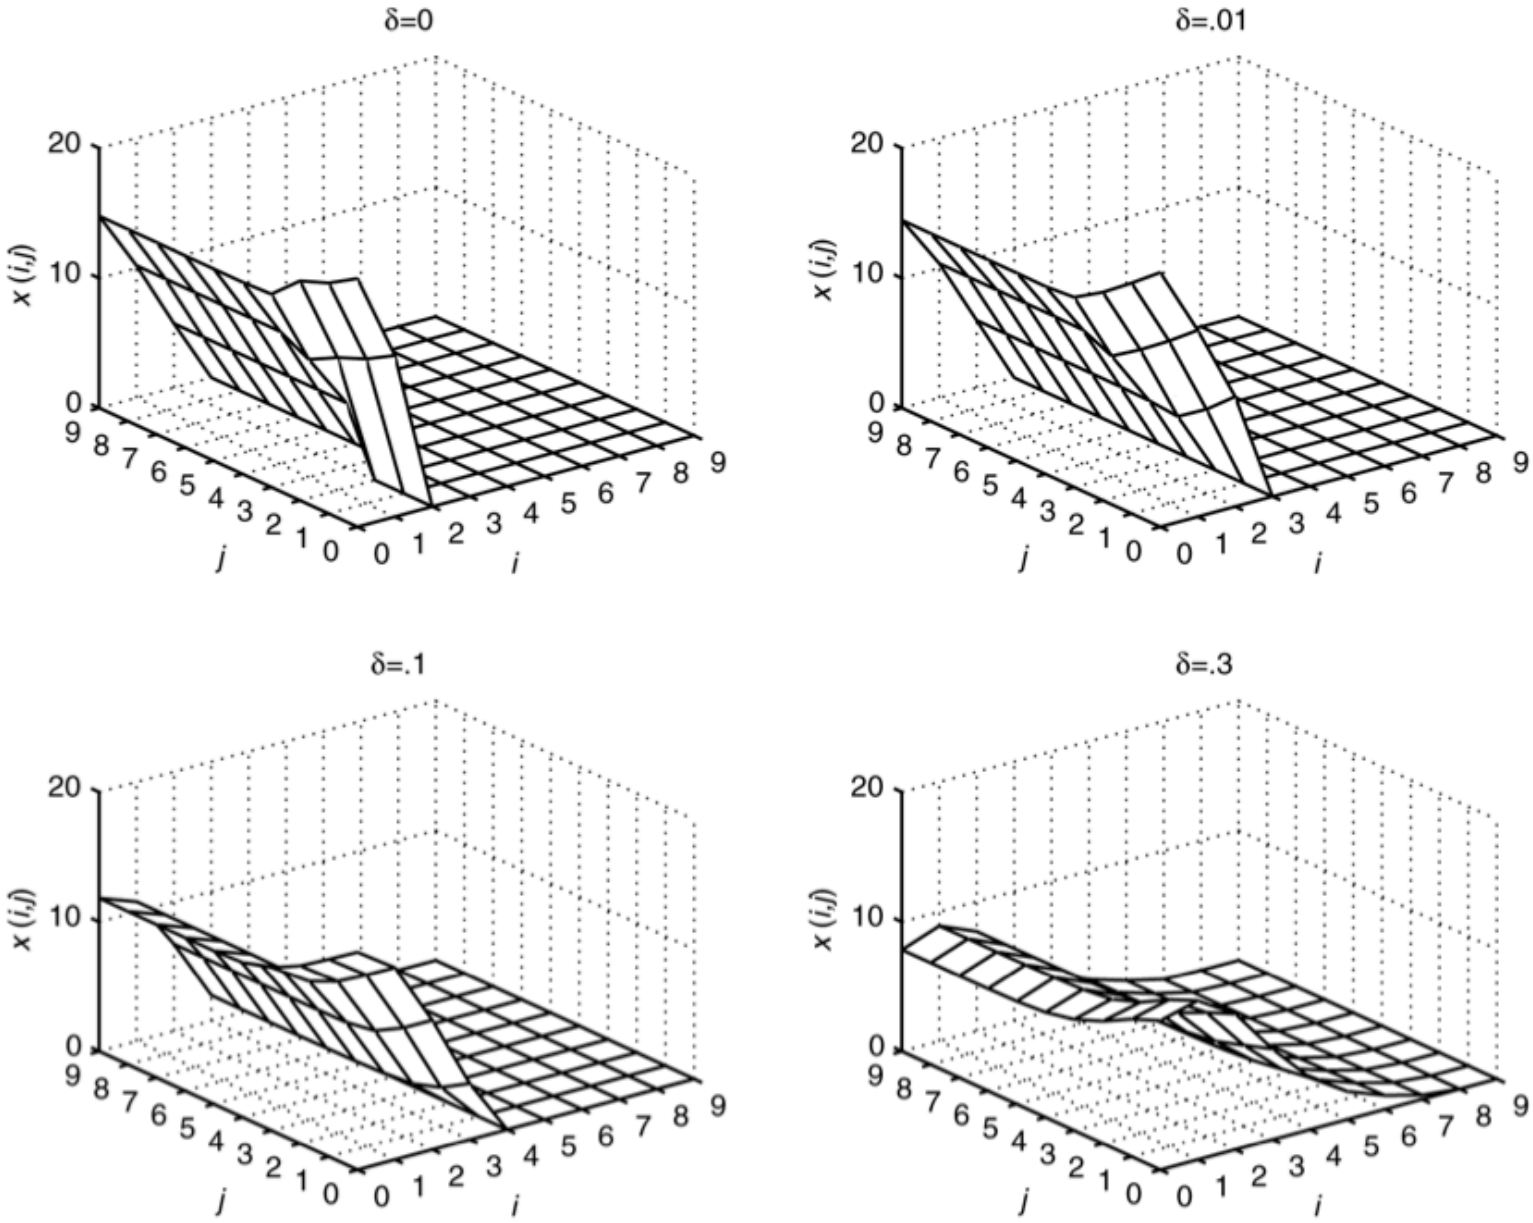
- **MPE Nash equilibrium:**
  - $x_i(\omega) = R_1(x_j(\omega)|\omega)$ for all $i, j = \{1, 2\}$.
  - $V_i^x(\omega) = \pi_i(\omega) + -x_i + \beta \sum_{\bar{q}_1'} W_i^x(\bar{q}_1'|\omega) \Pr(\bar{q}_i'|\bar{q}_i, x_i)$ for all $i, j = \{1, 2\}$.

Solution algorithm

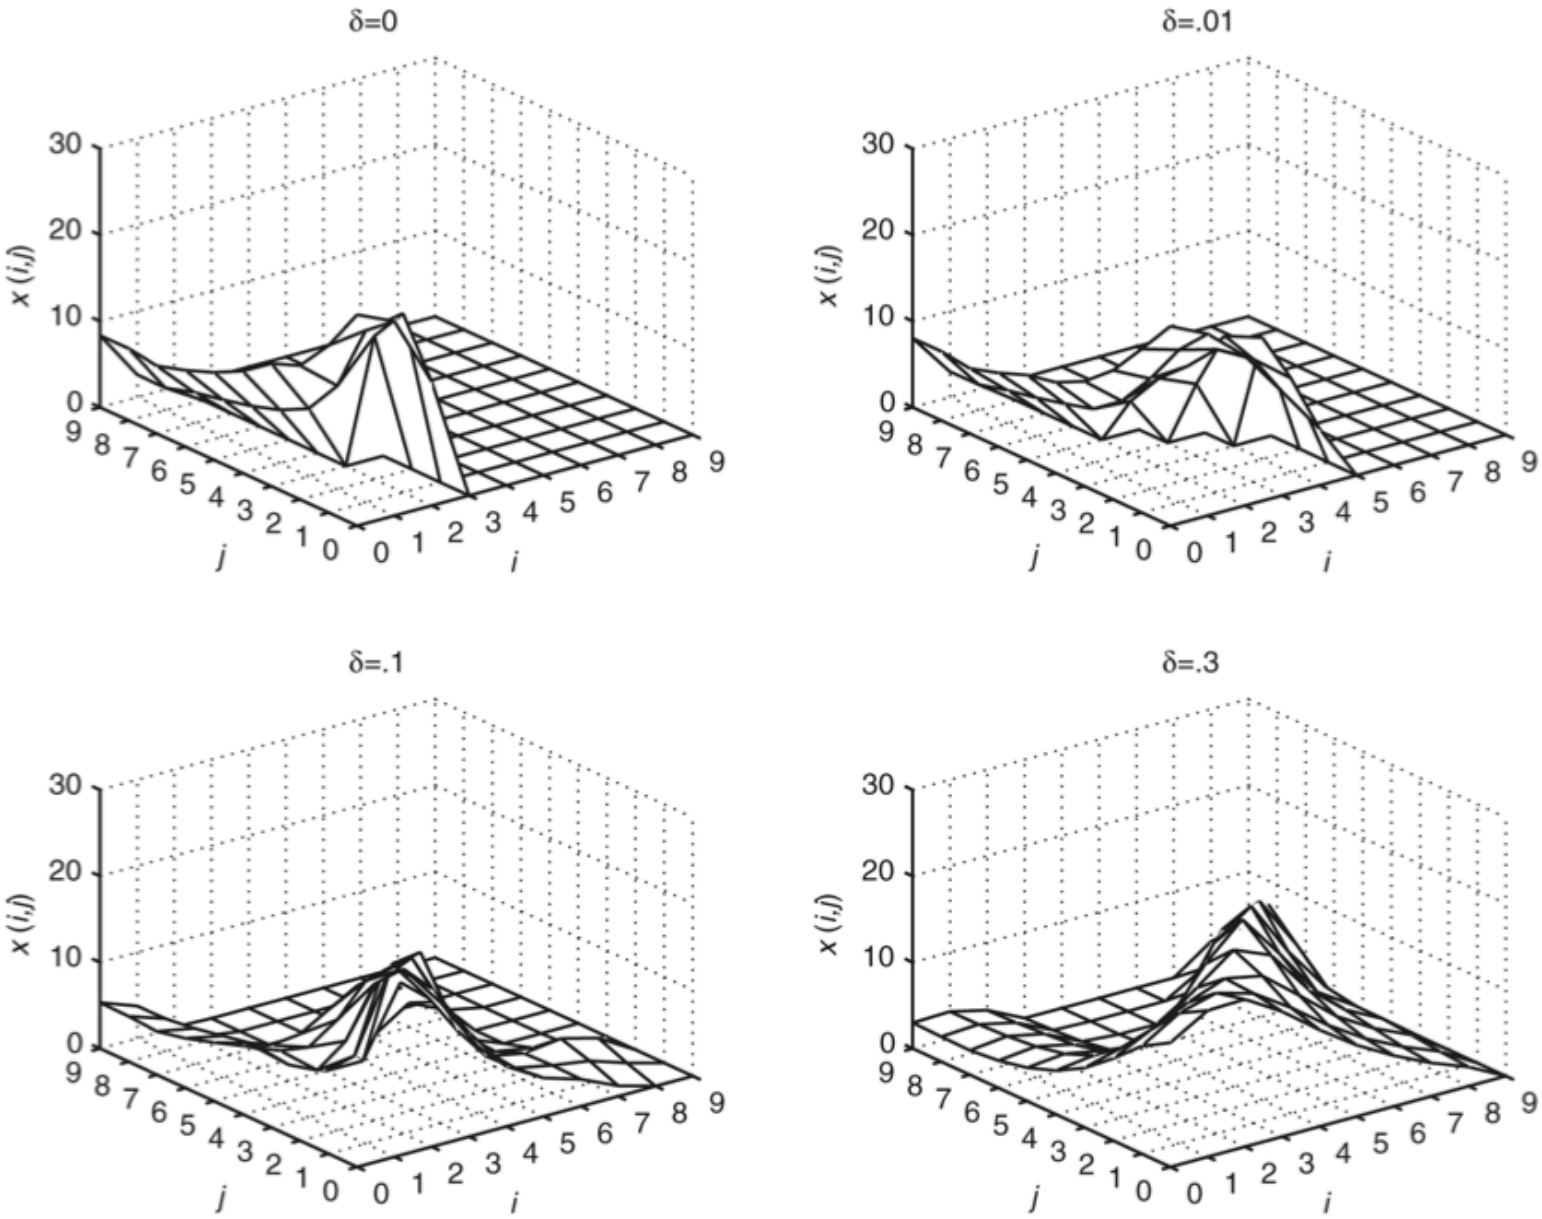
1. State space: $\bar{q}_i \in \{\bar{q}_0, \ldots, \bar{q}_{M-1}\}$ for $i = 1, 2$
2. Solve stage game profits for all $\omega = (\bar{q}_1, \bar{q}_2)$
3. Starting values: $x_i^1(\omega)$ and $V_i^1(\omega)$ for all ω for $i = 1, 2$ (imposing symmetry)
4. Iteration k:
   - Calculate continuation value:
   $$W_1^{k-1}(\bar{q}_1) = \sum_{\bar{q}_2'} V_1^{k-1}(\bar{q}_1, \bar{q}_2) \Pr(\bar{q}_2' | \bar{q}_2, x_2^{k-1})$$
   - Update $x_1(\omega)$ (see previous slide):
   $$x_1^k(\omega) = R_1(x^{k-1} | \omega)$$
   - Update $V_i(\omega)$:
   $$V_1^k(\omega) = \pi_1(\omega) - x_1^k(\omega) + \beta \sum_{\bar{q}_1'} W_1^{k-1}(\bar{q}_1' | \omega) \Pr(\bar{q}_1' | \bar{q}_1, x_1^k(\omega))$$
5. Repeat step (4) until: $||x_i^k(\omega) - x_i^{k-1}(\omega)|| < \epsilon_1$ and $||V_i^k(\omega) = V_i^{k-1}(\omega)|| < \epsilon_2$.

- **Simulations:**
  - Vary the depreciation rate δ, holding fixed the other parameters. Why?
  - δ measures the degree of investment irreversibility. Study the role of "credibility" on long-run industry distribution.
- **Simulation results**
  1. Quantity game
     - ★ If $\delta = 0$: Firms invest at constant rate until they reach Cournot: $x_1 = 0$ if $\bar{q}_1 > 3$.
     - ★ If $\delta > 0$: Firms over-invest (again +/- constant) to maintain the Cournot quantity in the long-run. As a result, firms hold idle capacity, the average long-run capacity levels are increasing in δ, and the industry structure tend to be symmetric.
  2. Pricing game
     - ★ Unlike in the Cournot model, investment functions depend crucially on competitors' capacity levels: Firms "give-up" quickly when opponent's level is larger than 3.
     - ★ When firms are symmetric, their investment levels increase: Firm engage in a "preemption" race to establish a higher capacity level.
     - ★ Investments are increasing in the depreciation fractor.
     - ★ Why? With price competition, small capacity firms are easily shut-down, and large capacity firms earn monopoly power on the residual demand segment.

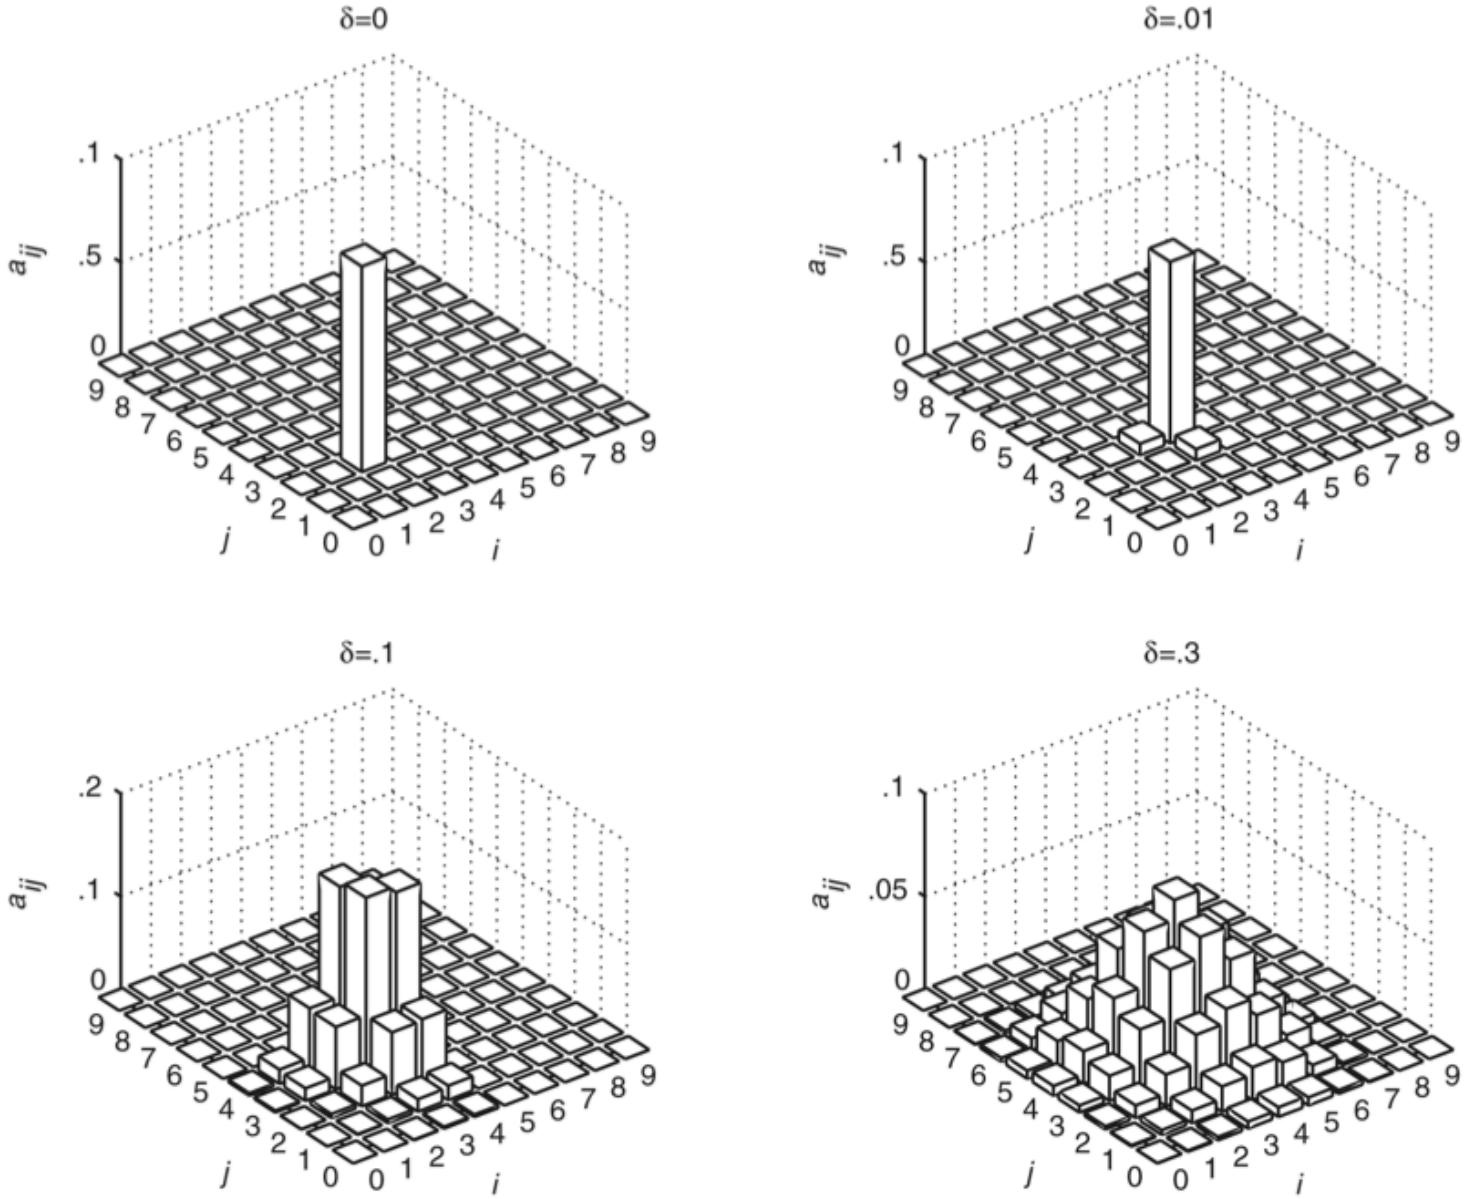
POLICY FUNCTION $x(i, j)$: QUANTITY COMPETITION



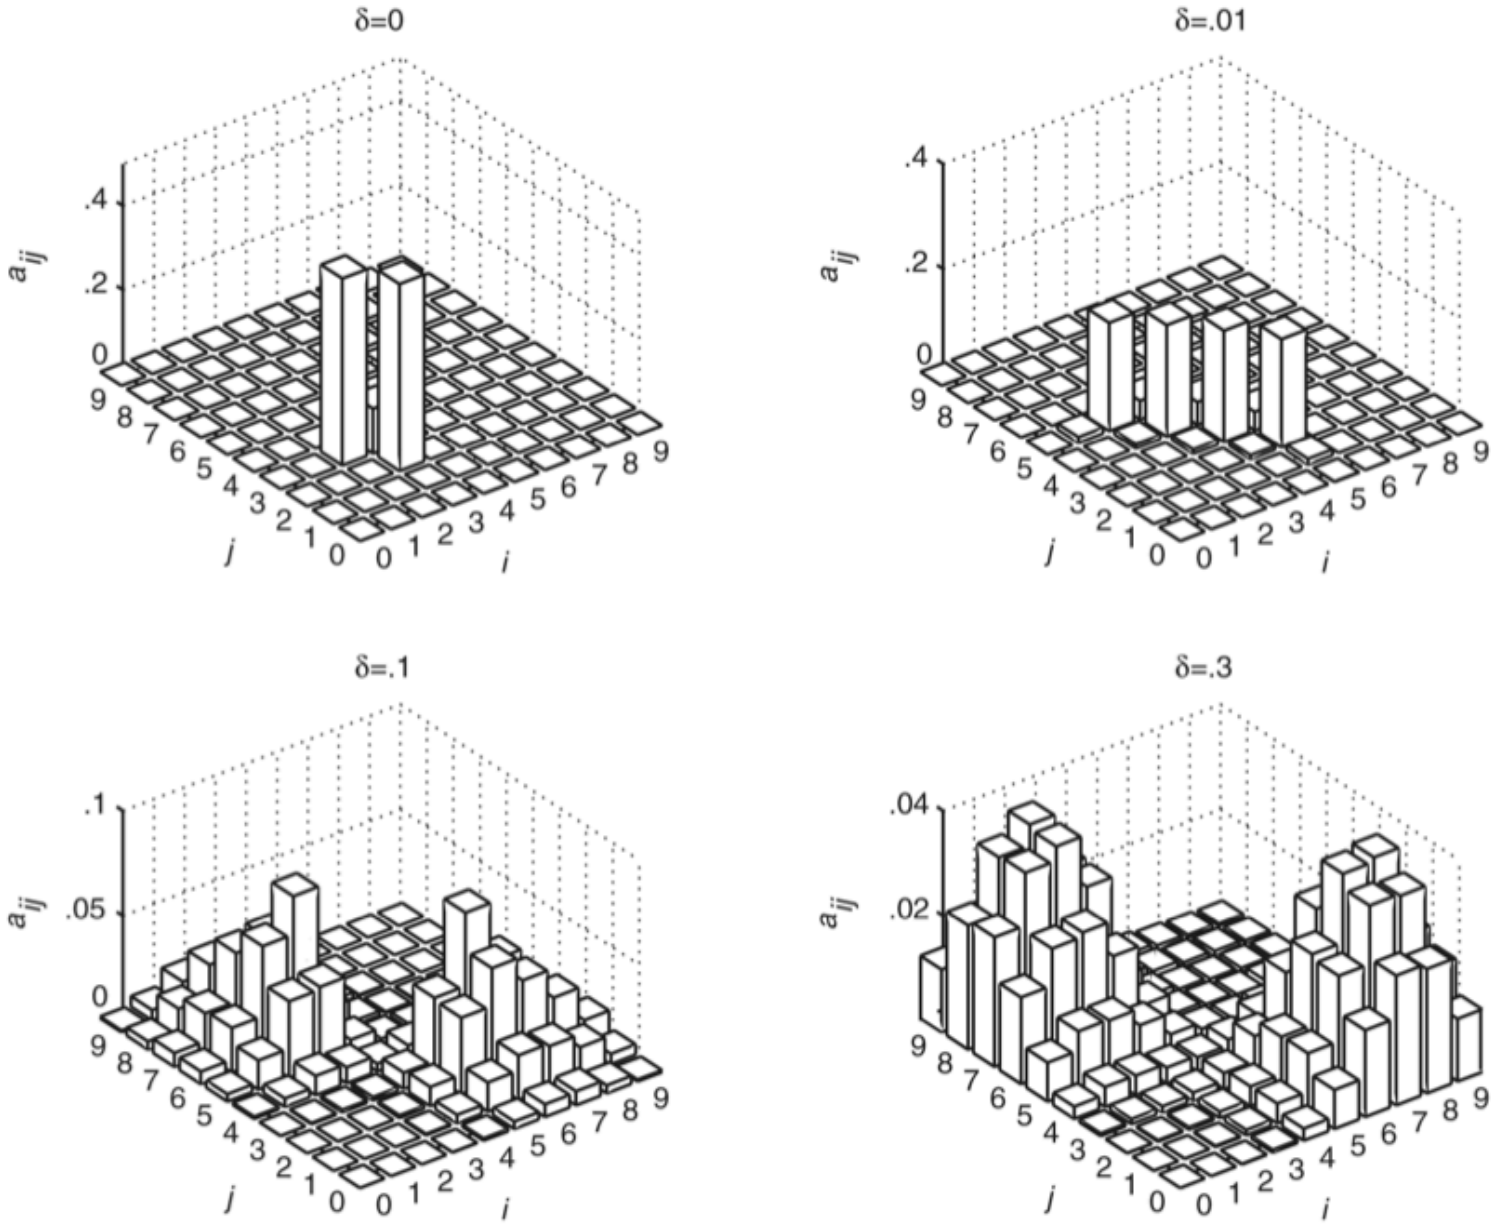
POLICY FUNCTION $x(i, j)$: PRICE COMPETITION



TRANSIENT DISTRIBUTION AFTER $T = 25$ PERIODS WITH INITIAL STATE $i_0 = j_0 = 1$: QUANTITY COMPETITION



TRANSIENT DISTRIBUTION AFTER $T = 25$ PERIODS WITH INITIAL STATE $i_0 = j_0 = 1$: PRICE COMPETITION



Markov perfect industry dynamics with many firms

- **Motivation:** Consider a game with N firms and M productivity levels
  - With $N = 20$ and $M = 40$, storying the policy functions of firms would require over 20 million gigabytes!
  - Even if it is storable, computing the transition of the industry is very costly, and would require an unreasonable level of rationality
  - Why? Firms forecast how their actions impact the behavior of all players in the market.

Markov perfect industry dynamics with many firms

- Available approaches to deal with "many firms" industry dynamics:
  - **Farias, Saure, and Weintraub (2012):** Value function approximation (next semester). Use basis functions to approximate the value function.
  - **Doraszelski and Judd (2012):** Continuous-time industry dynamics. Why? The probability that two firms choose actions (investment, entry, etc) simultaneously is zero. As a result firms believe that tomorrow's market structure is unchanged (not need to compute $Q(\omega'|\omega, x)$)
  - **Hopenhayn (1992):** Atomistic firms. When $N \to \infty$, the industry state is constant over time.
  - **Weintraub et al. (2008):** Obvilious equilibrium (today). N is finite, but firms choose their strategies assuming that the industry state is constant (invariant distribution). Similar to Hopenhayn, but rely on an approximation, rather than a limiting assumption. More suitable for empirical work.
  - **Ifrach and Weintraub (2017):** Moment-based Markov-perfect equilibrium. Firms choose their strategies using a (small) set of moments describing the state of the industry.

Oblivious equilibrium (Weintraub et al. (2008))

Model and notation

- Firms are differentiated in terms of quality: $x_{it} \in \{1, \ldots, K\}$
- Number of incumbents: n_t
- Industry state:

$$s_t = \left\{ \sum_{i=1}^{n_t} 1(x_{it} = 1), \sum_{i=1}^{n_t} 1(x_{it} = 2), \ldots, \sum_{i=1}^{n_t} 1(x_{it} = K) \right\} = \{s_t(1), \ldots, s_t(K)\}$$

$$s_{-i,t}(x) = \begin{cases} s_t(x) - 1 & \text{If } x = x_{it} \\ s_t(x) & \text{If } x \neq x_{it} \end{cases}$$

- Demand:

$$D(p_t | x_{it}, s_{-i,t}) = M \times \frac{\exp(g(x_{it}) - \alpha p_{it})}{1 + \exp(g(x_{it}) - \alpha p_{it}) + \sum_{k=1}^{K} s_{-i,t}(x_{it}) \exp(g(k) - \alpha p_{kt})}$$

  where M is the market size.

- Bertrand-Nash profit: $\pi(x_{it}, s_{-i,t})$

Actions and Timing

- Dynamic actions:
  - *Exit:* Private-value sell-off value ϕ_{it} (absorbing state)
  - *Entry:* Potential entrant can enter at quality level x^e by paying κ (common), or stay out (forever)
  - *Investment:* Incumbents choose to invest ι_{it} (marginal cost d) to improve the quality of their products

$$x_{i,t+1} = \max\left(0, x_{it} + w(\iota_{it}, \xi_{i,t+1})\right)$$

  where $\xi_{i,t+1}$ is a random-shock observed after investment is made.
- Timing of the game:
  1. Incumbents observe ϕ_{it} and make exit and investment decision
  2. Number of entrants is realized and entrants pay κ
  3. Incumbents compete in the spot market and profits are realized
  4. Exiting firms leave the market
  5. Investment outcomes are realized, new entrants enter, and industry state is updated: s_{t+1}

Assumptions

- Assumptions:
  - ϕ_{it} and ξ_{it} are IID
  - The investment return function, $w(\iota, \xi)$, is nondecreasing in ι
  - The transitions generated by $w(\iota, \xi)$ generate a unique investment choice
  - There exists a large number of potential entrants who play a mixed-strategy to enter the market.
- The last assumption implies that the distribution of entrants is well approximated by a Poisson distribution with mean $\lambda(s_t)$. In equilibrium $\lambda(s_t)$ is determined by a zero profit condition.
- Markov strategies:
  - Investment strategy: $\iota(x_{it}, s_{-i,t})$
  - Exit strategy: $\phi_{it} > \rho(x_{it}, s_{-i,t})$
  - Incumbents: $\mu = (\iota, \rho)$
  - Entry rate: $\lambda(s_t)$

Markov Perfect Equilibrium definition

- Incumbents value function given rivals' strategy (μ, λ):

$$V(x, s|\mu', \mu, \lambda) = \pi(x, s) + \max\left\{\phi_i, \max_{\iota} -d\iota + E_{\xi', s'}\left[V(x', s'|\mu, \lambda)|, \iota, x, s\right]\right\}$$

$$s.t. \quad x' = \max(0, x + w(\iota, \xi'))$$

$$\rho'(x, s) = -d\iota'(x, s) + E_{\xi', s'}\left[V(x', s'|\mu', \mu, \lambda)|, \iota'(x, s), x, s\right] \quad \text{[Exit cutoff]}$$

  where $\mu' = (\iota', \rho')$ denotes the optimal policy function given (μ, λ)

- Markov-perfect equilibrium:
  - Incumbent's strategies μ best-response to (μ, λ)
  - At each state, entrants earn zero profits or the entry rate is zero:

$$\beta E_{s'}\left[V(x^e, s'|\mu, \lambda)|s\right] - \kappa \leq 0 \forall s$$

$$\sum_s \lambda(s)\left(\beta E_{s'}\left[V(x^e, s'|\mu, \lambda)|s\right] - \kappa\right) = 0, \quad \lambda(s) \geq 0$$

- Large market observations:
  - There are no aggregate shocks, and players' actions are IID (conditional on s)
  - When the number of incumbents is large (i.e. m is large), simultaneous changes in individual firm quality levels can average out such that the industry structure remains +/- constant.

Oblivious strategies

- **Oblivious strategies:** Firms assume that the industry structure $s_{-i,t}$ is constant at its long-turn average, and condition their investment and exit strategies solely on their own quality level.

$$\text{Oblivious strategies: } \iota(x), \rho(x), \text{ and } \lambda.$$

- **Note:** If all firms use the same strategy (in equilibrium), the quality of each firm evolves according to independent transient Markov chains.
  - Let $P^1_\mu(x_t, x_{t+1}) = \Pr(x_{t+1}|x_t, \mu)$ denotes the one-period Markov process for quality (derived from strategy μ)
  - The two-period transition probability is:

$$P^2_\mu(x, y) = \sum_{x'} P^1_\mu(x, x') \Pr(x_{t+2} = y | x_{t+1} = x', \mu)$$

  - Similarly $P^k_\mu(x, x')$ is the k-period transition derived from μ
- **Main Implication:** The expected number of firms in state x if firms use strategy μ, λ is:

$$\lim_{t\to\infty} \tilde{s}_t(x) = \lambda \sum_{k=0}^{\infty} P^k_\mu(x|x^e) \equiv \tilde{s}_{\mu,\lambda}(x)$$

Oblivious equilibrium definition

- Given $(\tilde{s}_{\mu,\lambda}, \mu, \lambda)$, an Oblivious value function is given by:

$$\begin{aligned}
\tilde{V}(x|\mu', \mu, \lambda) &= \pi(x, \tilde{s}_{\mu,\lambda}) + \max\left\{\phi_i, \max_{\iota} -d\iota + E_{\xi'}\left[\tilde{V}(x'|\mu, \lambda)|, \iota, x\right]\right\} \\
&\quad s.t. \quad x' = \max(0, x + w(\iota, \xi')) \\
\Rightarrow &\quad \rho'(x) = -d\iota'(x) + E_{\xi'}\left[\tilde{V}(x'|\mu', \mu, \lambda)|, \iota'(x), x\right] \quad \text{[Exit cutoff]}
\end{aligned}$$

- Note that $\tilde{s}$ enters π through the demand and pricing game:

$$D(p_t|x_{it}, \tilde{s}) = M \times \frac{\exp(g(x_{it}) - \alpha p_{it})}{1 + \sum_{k=1}^{K} \tilde{s}(x_{it}) \exp(g(k) - \alpha p_{kt})}$$

- An **oblivious equilibrium** is a pair of strategy such that:
  - Incumbent strategies are optimal given (μ, λ) and $\tilde{s}_{\mu,\lambda}$
  - Entrant's earn zero profits:

$$\lambda(\beta \tilde{V}(x^e|\mu, \lambda) - \kappa) = 0$$

  - Beliefs about $\tilde{s}$ are consistent with (μ, λ):

$$\lambda \sum_{k=0}^{\infty} P_{\mu}^{k}(x|x^e) \equiv \tilde{s}_{\mu,\lambda}(x) \quad \forall x$$

Oblivious equilibrium definition (continued)

- Implications:
  - Under OE, the number of firms in each quality level is distributed according to a Poisson distribution that is independent across x.
  - When market size goes to infinity, the industry structure under MPE becomes approximately constant and converges to $\tilde{s}$.
- The previous results suggest that OE can approximate well MPE for large markets, but this is not quite true.
- *Light-tail assumption:* High-levels of concentration are unlikely to occur in the long-run as $M \to \infty$.
- When this assumption is satisfied, OE and MPE are approximately equal in large markets.
- In practice... this assumption cannot be verified easily since we cannot compute the MPE.
- Bottom line: (i) understand your model, and (ii) use OE as a "behavioral" assumption rather than strictly speaking as an approximation to MPE.

Computing OE

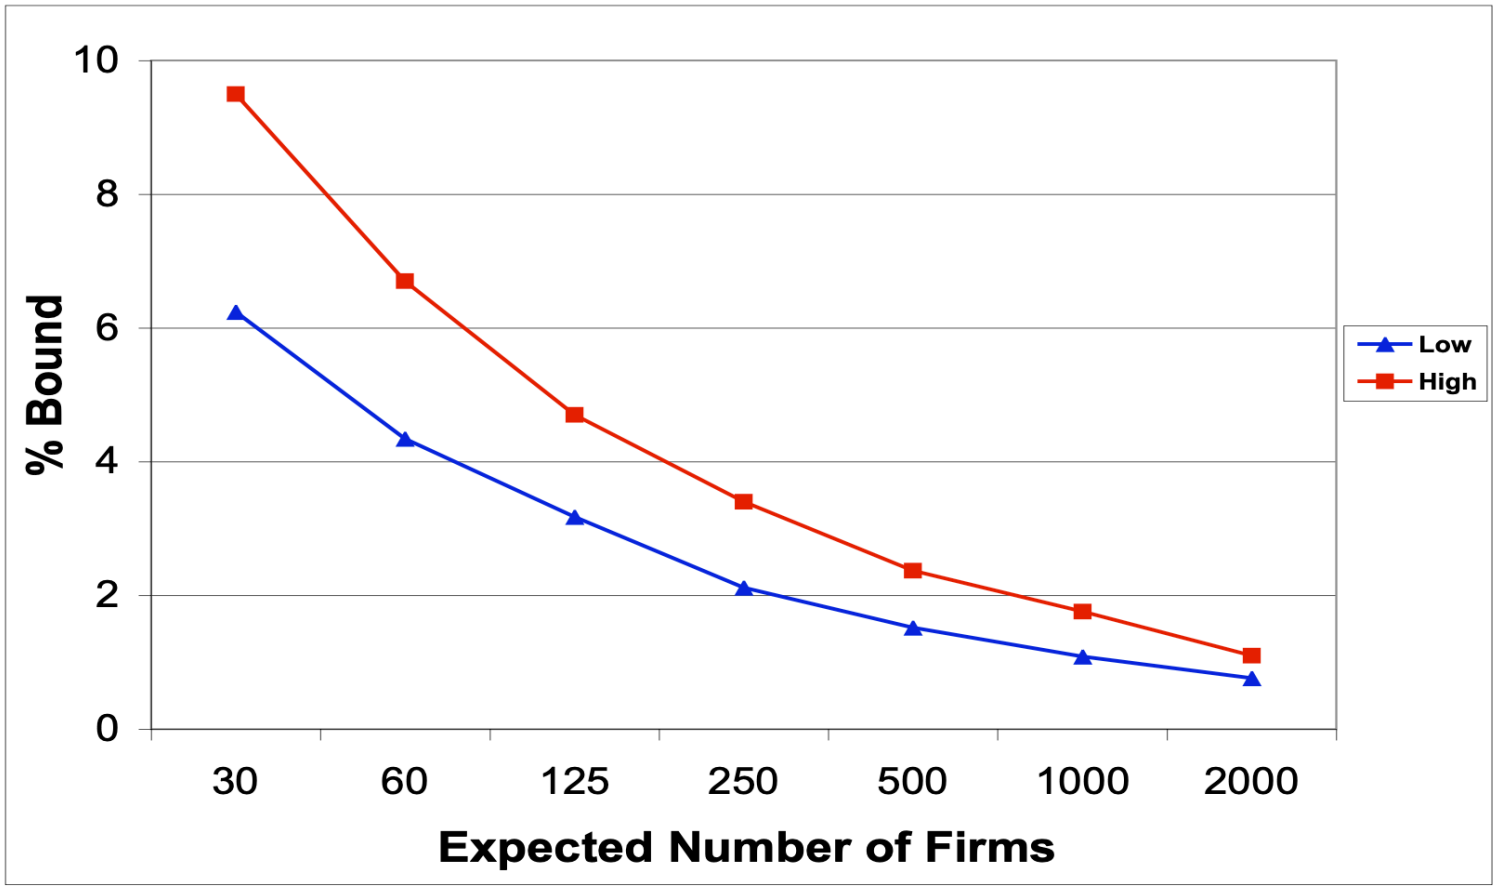
- The equilibrium entry rate lies between: $\underline{\lambda} = 0$, and $\overline{\lambda} = \frac{1}{\kappa}\left(\frac{\sup_{x,s}\pi(x,s)}{1-\beta} + \bar{\phi}\right)$

Algorithm 1 Oblivious Equilibrium Solver

1: $\underline{\lambda} := 0$; $\overline{\lambda} := \frac{1}{\kappa}\left(\frac{\sup_{x,s}\pi(x,s)}{1-\beta} + \bar{\phi}\right)$
2: $\mu(x) := 0$ for all x
3: $n := 0$
4: **repeat**
5: $\quad\lambda := (\underline{\lambda} + \overline{\lambda})/2$
6: $\quad$**repeat**
7: $\quad\quad$Choose $\mu^* \in \tilde{\mathcal{M}}$ to maximize $\tilde{V}(x|\mu^*, \mu, \lambda)$ simultaneously for all $x \in \mathbb{N}$
8: $\quad\quad\Delta := \|\mu^* - \mu\|_\infty$; $\quad n := n + 1$
9: $\quad\quad\mu := \mu + (\mu^* - \mu)/(n^\gamma + N)$
10: $\quad$**until** $\Delta \leq \epsilon_1$
11: $\quad$**if** $\beta\tilde{V}(x^e|\mu, \lambda) - \kappa \geq 0$ **then**
12: $\quad\quad\underline{\lambda} := \lambda$
13: $\quad$**else**
14: $\quad\quad\overline{\lambda} := \lambda$
15: $\quad$**end if**
16: **until** $|\beta\tilde{V}(x^e|\mu, \lambda) - \kappa| \leq \epsilon_2$

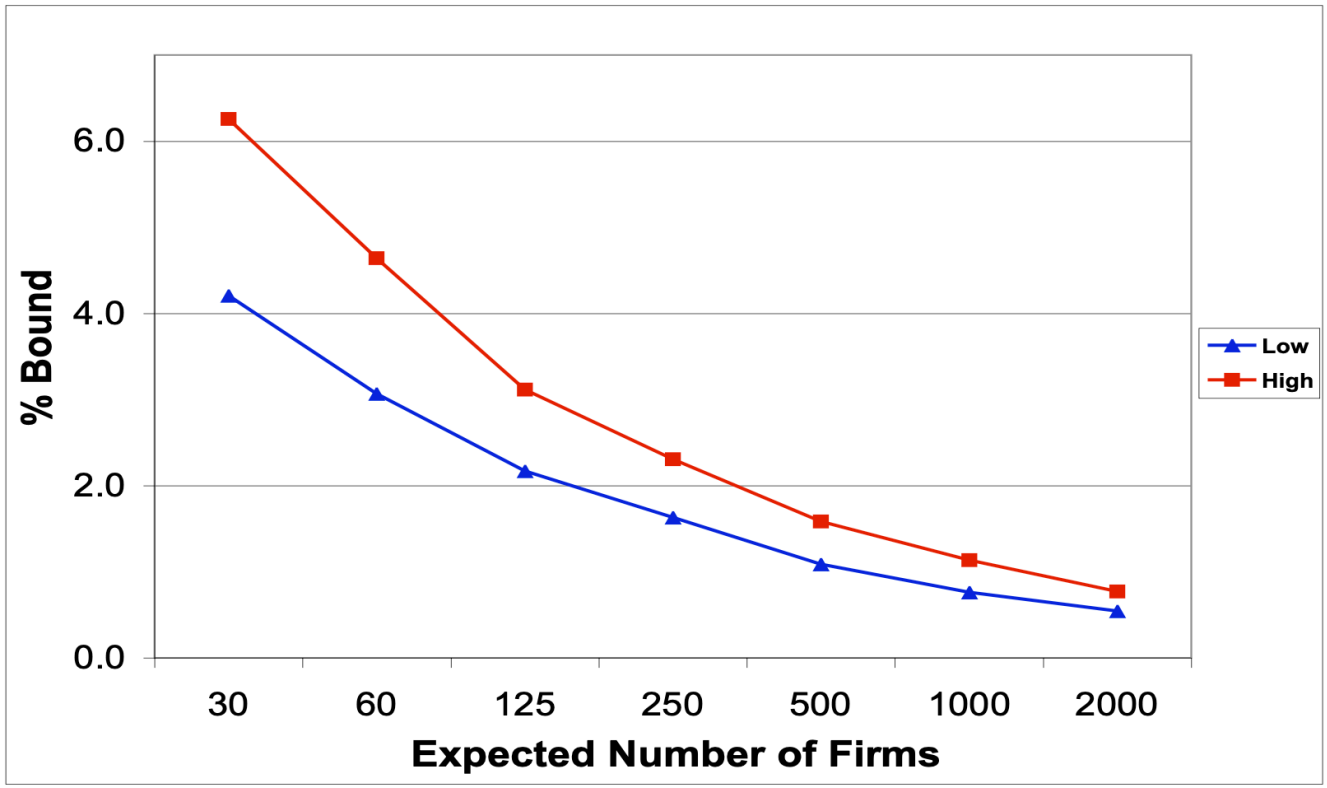
- Where γ and N are a smoothing parameter. Authors suggest: $\gamma = 2/3$ and $N = 0$.
- Source: Weintraub, Benkard, and van Roy (2010).
- The paper provides upper bounds on the approximation errors of using OE that do not require solving the MPE.

- Error bound in model with Poisson entry process and high or low levels of vertical differentiation



- Bottom line: Light-tail assumption is not satisfied when vertical/quality differentiation is important.

- Part of the approximation error is due to the fact that the number of entrants has a large variance (due to Poisson assumption)
- This next simulation exercise replaces the entry process by a deterministic entry:



- The approximation error is less than 2% when n is above 350 (compared to 1000+ previously)

- The model can be extended to incorporate aggregate shocks: Weintraub et al. (2010)
  - Aggregate shock (first-order markov): z_t
  - Extended strategies function of weighted averages of past shocks:

$$\iota(x_{it}, z_t, g_1(z_t, g_{t-1}), g_2(z_t, g_{t-2}), \ldots, g_N(z_t, g_{t-N})) = \iota(x_{it}, w_{it})$$

  - Long-run market structure: $\tilde{s}_{\mu,\lambda}(w_t)$

Algorithm 2 Extended Oblivious Equilibrium Solver

$\lambda(w) := \tilde{\lambda}$, for all w
$\mu(x, w) := \tilde{\mu}(x)$, for all x, w
$n := 0$
repeat
  Compute $\tilde{s}_{\mu,\lambda}(w)$, for all w
  Choose $\mu^* \in \tilde{\mathcal{M}}_z$ to maximize $\tilde{V}(x, w|\mu^*, \mu, \lambda)$ simultaneously for all x, w
  for all w **do**
    $\lambda^*(w) := \lambda(w) \left(\beta E\left[\tilde{V}(x^e, w_{t+1}|\mu^*, \mu, \lambda)\middle| w_t = w\right]/\kappa\right)$
  end for
  $\Delta_1 := \| \mu - \mu^* \|_\infty$, $\Delta_2 := \|\lambda - \lambda^*\|_\infty$
  $n := n + 1$
  $\mu := \mu + (\mu^* - \mu)/(n^{\gamma_1} + N_1)$
  $\lambda := \lambda + (\lambda^* - \lambda)/(n^{\gamma_2} + N_2)$
until $\Delta_1 \leq \epsilon_1$ and $\Delta_2 \leq \epsilon_2$

- Equilibrium concept is widely used when studying industry dynamics in large markets
  - Matching markets: Brancaccio, Kalouptsidi, and Papageorgiou (2020)
  - International trade: Chen and Xu (2020)

Limitation of oblivious equilibrium concept

- Poor approximation to MPE when the industry is likely to evolve to concentrate states
- Source of market dominance: Quality differentiation, R&D, Advertising, etc.
- Key: (i) ability for "productive" firms to grow in size, and (ii) increasing return to scale (cost or quality)

Example: Advertising and brand goodwill

- Demand given firm i's goodwill (x_{it}):

$$D(p_t|x_{it}, s) \propto x_{it}^{\alpha_1}(Y - p_{it})^{\alpha_2}$$

- Goodwill technology: Increasing return ($\alpha_1 > 1$), Constant return ($\alpha_1 = 1$), Decreasing return ($\alpha_1 < 1$)
- Long-run distribution of goodwill (x_{it}):

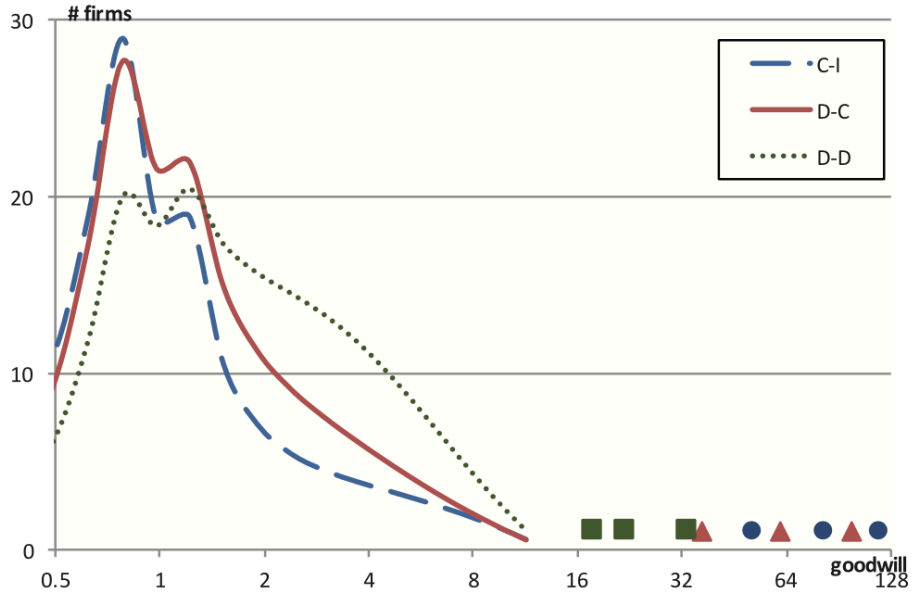


FIGURE 1

Long-run average size distribution (un-normalized) of firms in log-scale. The curve lines represent fringe firms for the different levels of returns to advertising. The squares represent dominant firms for D-D, the triangles for D-C, and the circles for C-I

Stylized facts

- Fact 1: Most concentrated markets have many small firms (e.g. Pareto distribution)
- Fact 2: The cereal example clearly illustrates that the advertising strategy of top brands affects the profit of others (including fringe). But the reverse is not true.
- Fact 3: The identity of top-brands is endogenous, and Fringe brands can become dominant if they are successful at innovating.
- Example: Chobani vs Danone and Yoplait

Moment-based equilibrium: Ifrach and Weintraub (2017)

- **Goal:** Tractable model of concentrated markets with many firms
- **Ingredients:**
  - *Exit:* Private-value sell-off value ϕ_{it} (absorbing state)
  - *Entry:* $N - n_t$ potential entrants face random entry-cost κ_{it}
  - *Investment:* Incumbents choose to invest ι_{it} (marginal cost d) to improve the quality of their products

$$x_{i,t+1} = \max\left(0, x_{it} + w(\iota_{it}, \xi_{i,t+1})\right)$$

    where $\xi_{i,t+1}$ is a random-shock observed after investment is made.
  - *Aggregate shock:* $z_t \sim p(z_t|z_{t-1})$ (e.g. demand)

Timing of the game

1. Incumbents observe ϕ_{it} and make exit and advertising decision
2. Potential entrants decide to enter or not after seeing κ_{it}
3. Incumbents compete in the spot market and profits are realized
4. Exiting firms leave the market
5. Investment outcomes are realized, new entrants enter, and industry state is updated: s_{t+1}

Dominant vs Fringe

- **Definition:** Firms with goodwill $x_{it} \geq x^d$ are "dominant" players, others are "fringe". The cutoff level is a "tuning" parameter.
- Industry state variables: $s = (d, f, z)$

$$f_t = \left\{ \sum_{i=1}^{N} 1(x_{it} = x_1), \sum_{i=1}^{N} 1(x_{it} = x_2), \ldots, \sum_{i=1}^{N} 1(x_{it} = x^d - \Delta) \right\}$$

$$d_t = \left\{ \sum_{i=1}^{N} 1(x_{it} = x^d), \ldots, \sum_{i=1}^{N} 1(x_{it} = x_K) \right\}$$

- **Alternative state-space:** $\hat{s}_t = (d_t, z_t, \theta_t)$, where $\theta_t = \text{Moments}(f_t)$ is a sufficient statistic of fringe distribution.
- **Example:** α-th un-normalized moment function

$$\theta_t^{\alpha} = \sum_{k=1}^{r(x^d)} x_k^{\alpha} f_t(x_k)$$

E.g: $\alpha = 1$ is the sum of goodwill among fringe incumbents.

Moment-based strategies

- Individual strategies:
  - Investment strategy: $\iota(x_{it}, \hat{s}_t)$
  - Exit strategy: $\phi_{it} > \rho(x_{it}, \hat{s}_t)$
  - Entry strategy: $\kappa_{it} < \lambda(\hat{s}_t)$
- Under MPE, the model generates a first-order Markov transition for the aggregate state: $P_{\mu,\lambda}(s'|s)$
- Since θ_t might not be a sufficient statistic for f_t, moment-based strategies might generate non-Markovian transitions for $\hat{s}$
- **Assumption:** $\hat{P}_{\mu,\lambda}(\hat{s}_{t+1}|\hat{s}_t) \approx \Pr_{\mu,\lambda}(\hat{s}_{t+1}|\hat{s}_t)$
- How to compute $\hat{P}_{\mu,\lambda}(\hat{s}_{t+1}|\hat{s}_t)$?

Case 1: Observed transition frequencies

- Given strategy (μ, λ), simulate action sequences from s_0 for $T \to \infty$
- Drop first T_0 periods and define the set of visited states (i.e. recurrent): $\tilde{s}_t \in \mathcal{R}_{\mu,\lambda}$. For $\tilde{s} \in \mathcal{R}_{\mu,\lambda}$:

$$\hat{P}_{\mu,\lambda}(\theta'|\hat{s}) = \lim_{T\to\infty} \sum_{t=t_0}^{T} \frac{1(\tilde{s}_t = \hat{s}, \theta_{t+1} = \theta')}{\sum_{t=t_0}^{T} 1(\hat{s}_t = \hat{s})}$$

- Similarly, the probability observed observing a Fringe transitioning to the dominant state:

$$\hat{P}_{\mu,\lambda}(f \to d|\hat{s}) = \lim_{T\to\infty} \sum_{t=t_0}^{T} \frac{1(\tilde{s}_t = \hat{s}, f \to d)}{\sum_{t=t_0}^{T} 1(\hat{s}_t = \hat{s})}$$

- Let $P_{\mu,\lambda}(x', d', z'|x, \tilde{s})$ denotes the "exact" Markov transition probability for the other state variables (i.e. MPE)

Case 1: Observed transition frequencies (continued)

- The joint transition probability of the states is:

$$\hat{P}_{\mu,\lambda}(x', \hat{s}' | \hat{s}, x) =$$

$$\begin{cases} P_{\mu,\lambda}(x', d', z' | x, \tilde{s}) \hat{P}_{\mu,\lambda}(\theta' | \hat{s}) \hat{P}_{\mu,\lambda}(f \to d | \hat{s}) & \text{If } \hat{s} \to \hat{s}' \text{ includes a new dominant player} \\ P_{\mu,\lambda}(x', d', z' | x, \tilde{s}) \hat{P}_{\mu,\lambda}(\theta' | \hat{s}) \left[1 - \hat{P}_{\mu,\lambda}(f \to d | \hat{s})\right] & \text{Otherwise.} \end{cases}$$

- Transitions from states outside of the recurrent class are set at arbitrary values.

Case 2: Deterministic transition

- Use the simulated transition to estimate a parametric transition function. E.g.:

$$\theta_{t+1} = g_0(d_t, z_t) + g_1(d_t, z_t)\theta_t \equiv G(\theta|d, z)$$

- Use $G(\theta|d, z)$ to predict transition of fringe state:

$$\hat{P}_{\mu,\lambda} = \begin{cases} 1 & \text{If } \theta' = G(\theta|d, z) \\ 0 & \text{If } \theta' \neq G(\theta|d, z) \end{cases}$$

- This is similar in spirit to Krusell and Smith (1998)

Next-step: How to compute $\pi(x, \tilde{s})$?

- Solution: Average profit conditional on θ
- Simulation procedure:
  - Candidate strategies: μ, λ
  - Simulate sequence of actions starting from s_0 for $T \to \infty$
  - Drop first T_0 periods and define the set of visited states (i.e. recurrent): $\tilde{s}_t \in \mathcal{R}_{\mu,\lambda}$.
  - For each $s_t = \tilde{s}$, calculate profit function: $\pi(x_{it}, d_t, f_t, z_t) = \pi(x_{it}, s_t)$
  - Average profits:

$$\hat{\pi}(x_{it}, d, z, \theta) = \lim_{T \to \infty} \sum_{t=t_0}^{T} \pi(x_{it}, d_t, f_t, z_t) \frac{1(d_t = d, z_t = z, \theta(f_t) = \theta_t)}{\sum_{t=t_0}^{T} 1(d_t = d, z_t = z, \theta(f_t) = \theta_t)}$$

Extensions: Smooth(er) transitions

- Interpolation:
  - Frequency estimator is sensitive to sampling noise.
  - Observed transitions from "nearby" states are informative:

$$\hat{P}_{\mu,\lambda}(\theta'|\hat{s}) = \lim_{T\to\infty} \sum_{t=t_0}^{T} \frac{K_h(\tilde{s}_t - \hat{s}, \theta_{t+1} - \theta')}{\sum_{t=t_0}^{T} K_h(\hat{s}_t - \hat{s})}$$

    where $K_h(\cdot)$ is a smooth Kernel.
  - The same interpolation strategy can be used to compute $\hat{\pi}(x_{it}, d, z, \theta)$.
  - You can also use regression/prediction models to predict both
- **Moment-based Markov Equilibrium (MME):**
  - Firms play Nash equilibrium strategies given perceived transitions $\hat{P}_{\mu,\lambda}(x', \hat{s}'|\hat{s}, x)$
  - Transitions are consistent with dominant and fringe strategies

Solution Algorithm

Algorithm 1 Equilibrium solver

Initialize with (μ, λ) and industry state $s_0 = (f_0, d_0, z_0)$ with corresponding $\hat{s}_0$;
$n := 1$, $\Delta := \epsilon + 1$;
while $\Delta > \epsilon$ **do**
  Simulate a T period sample path $\{(f_t, d_t, z_t)\}_{t=1}^{T}$ with corresponding $\{\hat{s}_t\}_{t=1}^{T}$ for large T;
  Calculate the observed frequencies of industry states $h(\hat{s}) := \frac{1}{T}\sum_{t=1}^{T} \mathbf{1}\{\hat{s}_t = \hat{s}\}$, for all $\hat{s} \in \hat{\mathcal{S}}$;
  Use the same sample path to compute the perceived transition kernel $\hat{\mathbf{P}}_{\mu',\mu,\lambda}[\cdot|x, \hat{s}]$, for all $x, \hat{s}$;
  Solve $\mu' := \underset{\tilde{\mu} \in \hat{\mathcal{M}}}{\operatorname{argmax}} \hat{V}(x, \hat{s}|\tilde{\mu}, \mu, \lambda)$, for all $(x, \hat{s}) \in \mathcal{X} \times \hat{\mathcal{S}}$;
  Let $\lambda'(\hat{s}) := \beta \mathrm{E}_{\mu,\lambda}[\hat{V}(x^e, \hat{s}_{t+1}|\mu', \mu, \lambda)|\hat{s}_t = \hat{s}]$, for all $\hat{s} \in \hat{\mathcal{S}}$;
  $\Delta := \max(\|\mu - \mu'\|_h, \|\lambda - \lambda'\|_h)$;
  $\mu := \mu + (\mu' - \mu)/(1 + n^{\sigma})$;
  $\lambda := \lambda + (\lambda' - \lambda)/(1 + n^{\sigma})$;
  $n := n + 1$;
end while;

- The stopping rule is based on a weighted-average norm. Higher weights on states that are more visited ($h(\tilde{s})$)
- Given transition $\hat{P}_{\mu,\lambda}(x', \hat{s}'|\hat{s}, x)$, the algorithm solves a single-agent dynamic programing problem.

Numerical example: Advertising as investment in goodwill

- Demand given firm i's goodwill (x_{it}):

$$D(p_t|x_{it}, s) \propto x_{it}^{\alpha_1}(Y - p_{it})^{\alpha_2}$$

- Return to advertising:

$$x_{i,t+1} = \begin{cases} \min(x_K, x_{it}(1+\Delta)) & \text{w.p. } \frac{\delta\psi(x)\iota_{it}}{1+\psi(x)\iota_{it}} \\ x_{it} & \text{w.p. } \frac{1-\delta'+(1-\delta)\psi(x)\iota_{it}}{1+\psi(x)\iota_{it}} \\ \min(x_1, x_{it}/(1+\Delta) & \text{w.p. } \frac{\delta'}{1+\psi(x)\iota_{it}} \end{cases}$$

- Three parameterizations:

$$(\alpha_1^f, \alpha_1^d) = \{\alpha_D, \alpha_C\} \times \{\alpha_D, \alpha_C, \alpha_I\}$$

and $(\alpha_D, \alpha_C, \alpha_I) = (0.85, 1, 1.1)$.

Other calibrated parameters (Beer industry)

Description	Value	Source
Number of firms (N)	200	This number is chosen to be greater than the maximal number of active firms in this period.
Maximal number of dominant firms ($\overline{D}$)	3	This is the actual number of dominant firms in the industry.
Depreciation of goodwill (ρ)	0.25	Roberts and Samuelson (1988) estimate this by .2 for the cigarette market.
Production cost per barrel (c)	\$120	Rojas (2008) estimated the markup to be about a third of the price, and the average price is \$165 per barrel in the period studied.
Average entry cost (exponential dist.)	35×10^6	Based on costs of new plants.
Sell-off value (exponential dist.)	7×10^6	Based on sales' price of used plants.
Fixed cost per period fringe (double for dominant)	10^6	We introduce this fixed cost in the single-period profit function.
Profit function parameters (Y and α_2)	200 and 1, resp.	Chosen to match the price elasticity (-0.5) for the average price, see (Tremblay and Tremblay, 2005, p. 23).

Long-run distribution of goodwill (x_{it})

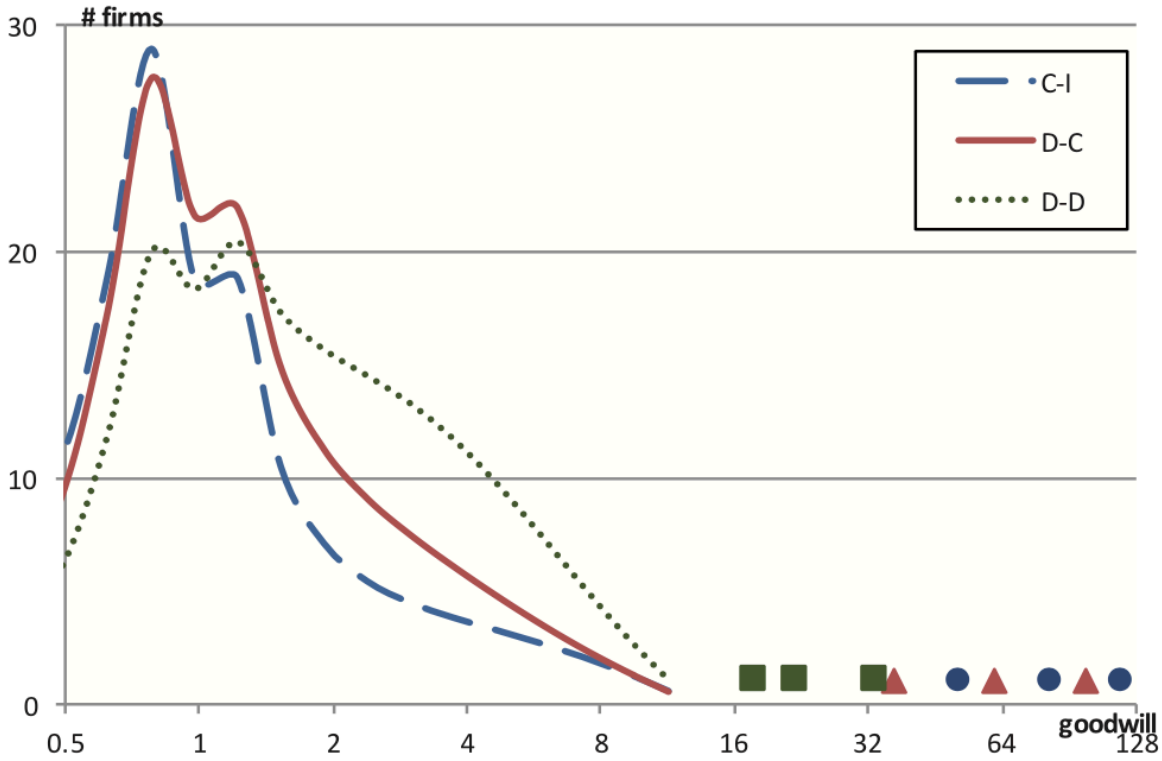


FIGURE 1

Long-run average size distribution (un-normalized) of firms in log-scale. The curve lines represent fringe firms for the different levels of returns to advertising. The squares represent dominant firms for D-D, the triangles for D-C, and the circles for C-I

Summary statistics

TABLE 3
Average industry statistics

Industry averages	CRA-IRA	DRA-CRA	DRA-DRA
Concentration ratio C_2	0.42	0.36	0.13
Concentration ratio C_3	0.53	0.44	0.17
First moment (normalized by # fringe)	1.09	1.22	1.74
Active fringe firms (#)	147	166	180
Active dominant firms (#)	3	3	2.3
Size incumbent fringe (goodwill)	1.1	1.4	2
Size dominant (goodwill)	83.4	64.9	24.1
Entrants per period (#)	16.5	12.3	7.6
Time in industry fringe	12.9	16	26.2
Time as dominant	1536	360	21

- What is the value-added of modeling the behavior of fringes?
- Assume that dominant firms are oblivious with respect to fringe state, $\theta_t = \hat{\theta}$ for all t, and compete by playing a MPE inestment game.
- **MPE vs MME:** The average size of dominant firms drops by 17%. And the average size of the smallest dominant firm falls around 50% in the MPE
- Why? Under MME, dominant players have incentives to deter entry and push down investment from fringe firms.

Benkard, L. (2004, July).
A dynamic analysis of the market for wide-bodied commercial aircraft.
Review of Economic Studies.

Besanko, D. and U. Doraszelski (2004).
Capacity dynamics and endogenous asymmetries in firm size.
Rand Journal of Economics.

Besanko, D., U. Doraszelski, Y. Kryukov, and M. Satterthwaite (2010, March).
Learning-by-doing, organizational forgetting, and industry dynamics.
Econometrica 78(2).

Brancaccio, G., M. Kalouptsidi, and T. Papageorgiou (2020).
Geography, transportation and endogenous trade costs.
Econometrica 88(2), 657–691.

Chen, Y. and D. Xu (2020, January).
A structural empirical model of r&d investment, firm heterogeneity, and industry evolution,.
Journal of International Economics.

Doraszelski, U. and K. Judd (2012, March).
Avoiding the curse of dimensionality in dynamic stochastic games.
Quantitive Economics, 53–93.
wp, Harvard University.

Doraszelski, U. and S. Markovich (2005).
Advertising dynamics and competitive advantage.
Forthcoming, Rand Journal of Economics.

Doraszelski, U. and A. Pakes (2007).
A framework for applied dynamic analysis in io.
In M. Armstrong and R. Porter (Eds.), *Handbook of Industrial Organization*, Volume 3, Chapter 30, pp. 1557–2440. Elsevier.

Doraszelski, U. and M. Satterthwaite (2005).
Foundations of markov-perfect industry dynamics: Existence, purification, and multiplicity.
mimeo, Harvard University.

Ericson, R. and A. Pakes (1995).
Markov-perfect industry dynamics: A framework for empirical work.
The Review of Economic Studies 62(1), 53–82.

Farias, V., D. Saure, and G. Y. Weintraub (2012, Summer).
An approximate dynamic programming approach to solving dynamic oligopoly models.
Rand Journal of Economics 43(2), 253–282.

Gowrisankaran, G. (1999).
A dynamic model of endogenous horizontal mergers.
Rand Journal of Economics 30, 56–83.

Hopenhayn, H. A. (1992).
Entry, exit, and firm dynamics in long run equilibrium.
Econometrica 60, 1127–1150.

Ifrach, B. and G. Y. Weintraub (2017).
A framework for dynamic oligopoly in concentrated industries.
Review of Economic Studies 84, 1106–1150.

Krusell, P. and J. Smith, Anthony A. (1998).
Income and wealth heterogeneity in the macroeconomy.
Journal of Political Economy 106(5), 867–896.

Pakes, A. and P. McGuire (1994).
Computing markov-perfect nash equilibria: Numerical implications of a dynamic differentiated product model.
The Rand Journal of Economics 25(4), 555–589.

Pakes, A. and P. McGuire (2001).
Stochastic algorithms, symmetric markov perfect equilibrium, and the 'curse' of dimensionality.
Econometrica 69(5), 1261–1281.

Weintraub, G. Y., L. Benkard, and B. van Roy (2008).
Markov perfect industry dynamics with many firms.
Econometrica 76(6), 1375–1411.

Weintraub, G. Y., L. Benkard, and B. van Roy (2010).
Computational methods for oblivious equilibrium.
Operations Research 58(4), 1247–1265.