An Introduction to Dynamic Games with Finite States

Ulrich Doraszelski

Agenda

- From dynamic programming to dynamic games.
- Value function iteration approach.
- Application: Quality ladder model without entry/exit.
- Markov-perfect industry dynamics.
- Existence, purification, and multiplicity of equilibrium.
- Application: Quality ladder model with entry/exit.
- Open questions.
- Alternative formulations: Continuous-time stochastic games with finite states.
- Application: Learning-by-doing.
- Computing all equilibria: Homotopy method.

From Dynamic Programming...

- Time is discrete. The horizon is infinite.
- The state space $\Omega = \{1, 2, \dots, L\}$ 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.

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

$$V_n(\omega) = \max_{x_n \in D_n(\omega)} \pi_n(\omega, x_n, X_{-n}(\omega)) + \beta \sum_{\omega_1'=1}^{L_1} \ldots \sum_{\omega_N'=1}^{L_N} V_n(\omega') \mathrm{Pr}(\omega' | \omega, x_n, X_{-n}(\omega)),$$

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

Value Function Iteration Approach

- Order the states in Ω and make initial guesses for the value $V_n(\omega)$ and policy $X_n(\omega)$ of each player $n = 1, \ldots, N$ in each state $\omega \in \Omega$.

- Proceed through the state space Ω in the prespecified order. In state $\omega \in \Omega$, given old guesses $V_n(\omega)$ and $X_n(\omega)$ compute new guesses $\widehat{V}_n(\omega)$ and $\widehat{X}_n(\omega)$ for each player $n = 1, \ldots, N$ as follows:

$$\widehat{X}_n(\omega) = \arg\max_{x_n \in D_n(\omega)} \pi_n(\omega, x_n, X_{-n}(\omega)) + \beta \sum_{\omega'=1}^{L} V_n(\omega') \Pr(\omega' | \omega, x_n, X_{-n}(\omega)),$$

$$\widehat{V}_n(\omega) = \pi_n\left(\omega, \widehat{X}_n(\omega), X_{-n}(\omega)\right) + \beta \sum_{\omega'=1}^{L} V_n(\omega') \Pr(\omega' | \omega, \widehat{X}_n(\omega), X_{-n}(\omega)).$$

  $\rightarrow$ Gauss-Jacobi scheme at each state $\omega \in \Omega$.

- Two ways to update $V_n(\omega)$ and $X_n(\omega)$:
  - Use a Gauss-Jacobi scheme: Replace old with new guesses after a complete pass through the state space (Pakes & McGuire 1994).
  - Use a block Gauss-Seidel scheme: Immediately after visiting a state replace old with new guesses for that state.

- Continue cycling through the state space until the changes in the value and policy functions are “small.”

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

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

where

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

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

  - $x_{-n}(\omega)$ 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

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

   and

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

   where

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

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.

Markov-Perfect Industry Dynamics

- Ericson, R. & Pakes, A. (1995) "Markov-Perfect Industry Dynamics: A Framework for Empirical Work."

- The EP model of dynamic competition in an oligopolistic industry is a special case of a dynamic game:
  - Entry, exit, and investment decisions.
  - Product market competition.

- It captures two key findings of the empirical literature:
  - Entry and exit occur simultaneously.
  - Heterogeneity among firms evolves endogenously.

- It has been widely applied in IO and other fields.

Markov-Perfect Industry Dynamics

- Advertising (Doraszelski & Markovich 2007).
- Capacity accumulation (Besanko & Doraszelski 2004, Chen 2009, Ryan 2012, Besanko, Doraszelski, Lu & Satterthwaite 2010a, 2010b, Wilson 2012).
- Collusion (Fershtman & Pakes 2000, 2005, de Roos 2004).
- Competitive convergence (Langohr 2003).
- Consumer learning (Ching 2010).
- Corporate reputation (Abito, Besanko & Diermeier 2012).
- Learning by doing (Benkard 2004, Besanko, Doraszelski, Kryukov & Satterthwaite 2010, Besanko, Doraszelski & Kryukov 2013).
- Mergers (Berry & Pakes 1993, Gowrisankaran 1999, Mermelstein, Nocke, Satterthwaite & Whinston 2013).
- Network effects (Jenkins, Liu, Matzkin & McFadden 2004, Markovich 2004, Markovich & Moenius 2005, Chen, Doraszelski & Harrington 2009).
- Productivity growth (Laincz 2005).
- R&D (Gowrisankaran & Town 1997, Auerswald 2001, Song 2011).
- Switching costs (Chen 2011).
- Technology adoption (Schivardi & Schneider 2005).
- International trade (Erdem & Tybout 2003).
- Finance (Goettler, Parlour & Rajan 2004).

Existence, Purification, and Multiplicity of Equilibrium

- Doraszelski, U. & Satterthwaite, M. (2010) "Computable Markov-Perfect Industry Dynamics."

- Questions:
  - Does a MPE exist in the EP model?
  - Is the MPE computationally tractable?
    * Pure strategies.
    * Symmetric and anonymous (exchangeable).
  - Is the MPE unique?

- Answers:
  - In the EP model a symmetric and anonymous MPE in pure strategies always exists under reasonable conditions.
  - The MPE is not necessarily unique.

Three Difficulties

- Randomization over discrete actions (entry/exit):
  - Introduce randomly drawn, privately-known setup costs/scrap values $\rightarrow$ the game of incomplete information has a MPE in cutoff entry/exit strategies.
- Randomization over continuous actions (investment):
  - Provide conditions on the model's primitives (UIC admissibility) such that a firm's optimal investment level is always unique $\rightarrow$ the MPE is in pure investment strategies.
  - Recent generalization: Escobar, J. (2013) "Equilibrium Analysis of Dynamic Models of Imperfect Competition."
- Symmetry and anonymity.
  - Provide conditions on the model's primitives $\rightarrow$ the MPE is symmetric and anonymous.

Quality Ladder Model with 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. (2012) "A Dynamic Quality Ladder Model with Entry and Exit: Exploring the Equilibrium Correspondence Using the Homotopy Method."

- Incumbent firms (i.e., active firms) and potential entrants (i.e., inactive firms).

- Two firms that can be either a potential entrant or an incumbent firm with potentially different product qualities

$$\omega = (\omega_1, \omega_2) \in \{ \underbrace{1, \ldots, L}_{\text{active firm}}, \underbrace{L+1}_{\text{inactive firm}} \}^2 = \Omega.$$

- Exit is a transition from state $\omega_n \neq L+1$ to state $\omega_n' = L+1$.

- Entry is a transition from state $\omega_n = L+1$ to state $\omega_n' = \omega^e \neq L+1$, where ω^e is an exogenously given initial product quality.

Quality Ladder Model with Entry/Exit

- Let $\xi_n(\omega) \in [0,1]$ be firm n's probability of remaining in (if $\omega_n \neq L+1$) or entering into (if $\omega_n = L+1$) the industry.

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

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

etc. If $\omega_n = L+1$, then

$$\Pr(\omega_n'|\omega_n,\xi_n) = \begin{cases} \xi_n & \text{if} \quad \omega_n' = \omega^e, \\ 1-\xi_n & \text{if} \quad \omega_n' = L+1. \end{cases}$$

Quality Ladder Model with Entry/Exit

- Firm n is assigned a random scrap value $\phi_n \sim F$ (if $\omega_n \neq L+1$) or a random setup cost $\phi_n^e \sim F^e$ (if $\omega_n = L+1$).
  - Scrap values/setup costs are private information.
  - Scrap values/setup costs are independent across firms and periods.

- Because scrap values and setup costs are private to a firm, its rivals perceive the firm *as if* it is mixing.

- In each period the timing is as follows:
  - Incumbent firms learn their scrap value and decide on exit and investment. Potential entrants learn their setup cost and decide on entry and investment.
  - Incumbent firms compete in the product market.
  - Exit and entry decisions are implemented.
  - The investment decisions of the remaining incumbents and new entrants are carried out and their uncertain outcomes are realized.

Incumbent Firm

- Bellman equation without entry/exit:

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

- Bellman equation with entry/exit:

$$V_n(\omega) = \max_{\xi_n \in [0,1], x_n \geq 0} \pi_n(\omega) + (1 - \xi_n) \mathrm{E} \left\{ \phi_n | \phi_n \geq F^{-1}(\xi_n) \right\}$$

$$+ \xi_n \left\{ -x_n + \beta \sum_{\omega_n'=1}^{L} W_n(\omega_n'; \omega_{-n}, \xi_{-n}(\omega), x_{-n}(\omega)) \mathrm{Pr}(\omega_n' | \omega_n, x_n, \xi_n = 1) \right\},$$

where

$$(1 - \xi_n) \mathrm{E} \left\{ \phi_n | \phi_n \geq F^{-1}(\xi_n) \right\} = \int_{\phi_n \geq F^{-1}(\xi_n)} \phi_n dF(\phi_n).$$

- An optimizing incumbent cares about the scrap value conditional on receiving it.

- Optimality condition:

$$\xi_n(\omega) = F \left(-x_n(\omega) + \beta \sum_{\omega_n'=1}^{L} W_n(\omega_n'; \omega_{-n}, \xi_{-n}(\omega), x_{-n}(\omega)) \mathrm{Pr}(\omega_n' | \omega_n, x_n(\omega), \xi_n = 1) \right)$$

Potential Entrant

- Potential entrants are short-lived.

- Bellman equation:

$$V_n(\omega) = \max_{\xi_n \in [0,1]} \xi_n \Bigg\{ -\mathsf{E}\left\{\phi_n^e | \phi_n^e \leq F^{e-1}(\xi_n)\right\}$$
$$+\beta \sum_{\omega_n'=1}^{L} W_n(\omega_n'; \omega_{-n}, \xi_{-n}(\omega), x_{-n}(\omega)) \Pr(\omega_n' | \omega_n, \xi_n = 1) \Bigg\},$$

  where

$$\xi_n \mathsf{E}\left\{\phi_n^e | \phi_n^e \leq F^{e-1}(\xi_n)\right\} = \int_{\phi_n^e \leq F^{e-1}(\xi_n)} \phi_n^e dF^e(\phi_n^e).$$

- An optimizing entrant cares about the setup cost conditional on paying it.

- Optimality condition:

$$\xi_n(\omega) = F^e \left(\beta \sum_{\omega_n'=1}^{L} W_n(\omega_n'; \omega_{-n}, \xi_{-n}(\omega), x_{-n}(\omega)) \Pr(\omega_n' | \omega_n, \xi_n = 1) \right).$$

Multiplicity

. . . we have experimented quite a bit with the core version of the algorithm, and we never found two sets of equilibrium policies for a given set of primitives (we frequently run the algorithm several times using different initial conditions or different orderings of points looking for other equilibria that might exist). We should emphasize here that the core version, and indeed most other versions that have been used, all use quite simple functional forms for the primitives of the problem, and multiplicity of equilibrium may well be more likely when more complicated functional forms are used. Of course, most applied work suffices with quite simple functional forms. (Pakes 2000, pp. 18–19)

Multiplicity: Quality Ladder Model without Entry/Exit

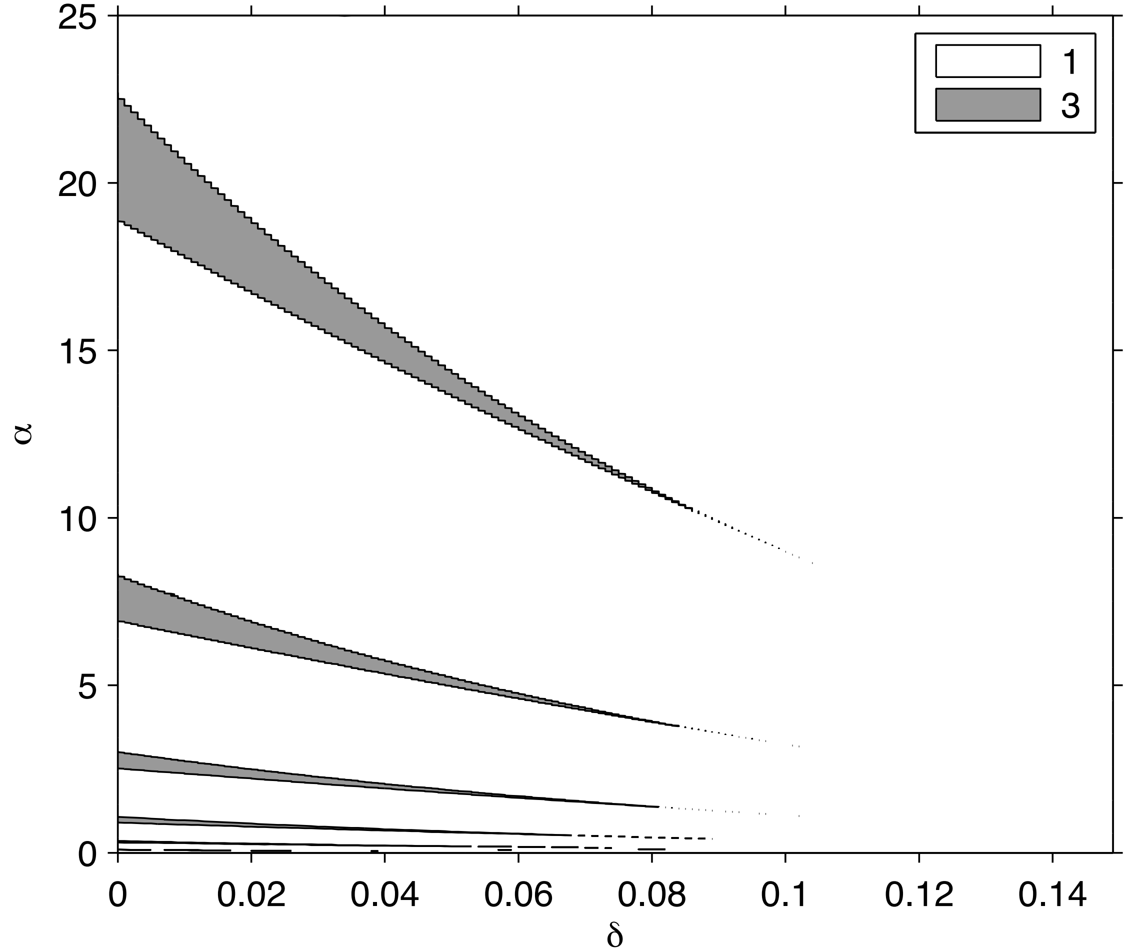

Number of equilibria in the Pakes & McGuire (1994) quality ladder model without entry/exit.
Source: Borkovsky, Doraszelski & Kryukov (2010).

Multiplicity: Quality Ladder Model with Entry/Exit

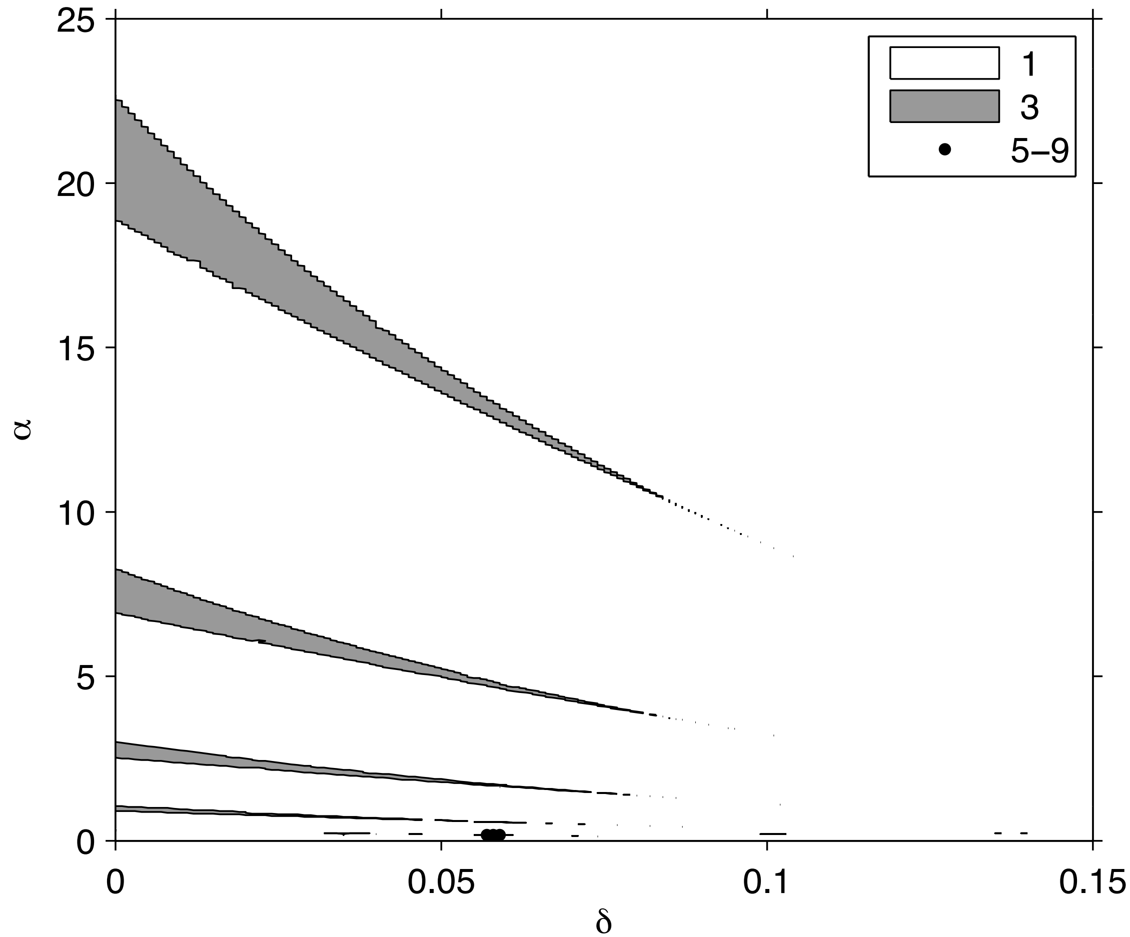

Number of equilibria in the Pakes & McGuire (1994) quality ladder model with entry/exit.
Source: Borkovsky, Doraszelski & Kryukov (2012).

Multiplicity

. . . I should note that virtually all Markov Perfect Models have multiple equilibria. . . (anonymous referee, 2013)

Open Questions: Computation

How can we compute equilibria faster?

- Pakes, A. & McGuire, P. (2001) "Stochastic Algorithms, Symmetric Markov Perfect Equilibrium, and the 'Curse' of Dimensionality."
- Ferris, M., Judd, K. & Schmedders, K. (2007) "Solving Dynamic Games with Newton's Method."
- Farias, V., Saure, D. & Weintraub, G. (2012) "An Approximate Dynamic Programming Algorithm to Solving Dynamic Oligopoly Models."
- Santos, C. (2012) "An Aggregation Method to Solve Dynamic Games."
- Arcidiacono, P., Bayer, P., Bugni, F. & James, J. (2011) "Sieve Value Function Iteration for Large State Space Dynamic Games."
- Aguirregabiria, V. and Vincentini, G. (2012) "Dynamic Spatial Competition Between Multi-Store Firms."
- Judd, K., Maliar, L. & Maliar, S. (2012) "Merging Simulation and Projection Approaches to Solve High-Dimensional Problems."

Open Questions: Computation

How can we compute more equilibria?

- Besanko, D., Doraszelski, U., Kryukov, Y. & Satterthwaite, M. (2010) "Learning-by-Doing, Organizational Forgetting, and Industry Dynamics."

- Borkovsky, R., Doraszelski, U. & Kryukov, Y. (2010) "A User's Guide to Solving Dynamic Stochastic Games Using the Homotopy Continuation Method."

Open Questions: Computation

Are other model formulations or equilibrium concepts computationally less burdensome?

- Doraszelski, U. & Judd, K. (2011) "Avoiding the Curse of Dimensionality in Dynamic Stochastic Games."
- Arcidiacono, P. Bayer, P. Blevins, J. & Ellickson (2012) "Estimation of Dynamic Discrete Choice Models in Continuous Time."
- Doraszelski, U. & Judd, K. (2007) "Dynamic Stochastic Games with Sequential State-to-State Transitions."
- Weintraub, G., Benkard, L. & Van Roy, B. (2008) "Markov Perfect Industry Dynamics with Many Firms."
- Weintraub, G., Benkard, L. & Van Roy, B. (2010) "Computational Methods for Oblivious Equilibrium."
- Ifrach, B. and Weintraub, G. (2012) "A Framework for Dynamic Oligopoly in Concentrated Industries."

Open Questions: Theory

What do we know about the general properties of the set of equilibria?

- Doraszelski, U. & Escobar, J. (2010) "A Theory of Regular Markov Perfect Equilibria in Dynamic Stochastic Games: Genericity, Stability, and Purification."

What types of behaviors can arise?

- Besanko, D., Doraszelski, U., Kryukov, Y. & Satterthwaite, M. (2010) "Learning-by-Doing, Organizational Forgetting, and Industry Dynamics."

- Judd, K. & Yeltekin, S. (2007) "Computing Equilibrium Value Sets for Dynamic Games with State Variables."

- Doraszelski, U. & Escobar, J. (2012) "Restricted Feedback in Long Term Relationships."

- Balbus, L., Reffett, K. & Wozny, L. (2010) "A Constructive Study of Markov Equilibria in Stochastic Games with Strategic Complementarities."

Open Questions: Applications

How can we deal with multiplicity in estimation?

- Aguirregabiria, V. & Mira, P. (2007) "Sequential Estimation of Dynamic Discrete Games."
- Bajari, P., Benkard, L. & Levin, J. (2007) "Estimating Dynamic Models of Imperfect Competition."
- Pakes, A., Ostrovsky, M. & Berry, S. (2007) "Simple Estimators for the Parameters of Discrete Dynamic Games (with Entry/Exit Examples)."
- Pesendorfer, M. & Schmidt-Dengler, P. (2008) "Asymptotic Least Squares Estimators for Dynamic Game."
- Judd, K. & Su, C. (2012) "Constrained Optimization Approaches to Estimation of Structural Models."

Open Questions: Applications

How can we deal with multiplicity in counterfactuals? How do players learn to play an equilibrium?

- Lee, R. & Pakes, A. (2009) "Multiple equilibria and selection by learning in an applied setting."

- Doraszelski, U. & Escobar, J. (2010) "A Theory of Regular Markov Perfect Equilibria in Dynamic Stochastic Games: Genericity, Stability, and Purification."

- Aguirregabiria, V. (2012) "A Method for Implementing Counterfactual Experiments in Models With Multiple Equilibria."

How can we deal with persistent asymmetric information?

- Fershtman, C. & Pakes, A. (2012) "Dynamic Games With Asymmetric Information: A Framework For Empirical Work."

- Bernhardt, D. & Taub, B. (2012) "Oligopoly Learning Dynamics."

Continuous-Time Stochastic Games with Finite States

- Doraszelski, U. & Judd, K. (2011) "Avoiding the Curse of Dimensionality in Dynamic Stochastic Games."

- Discrete-time, finite-state stochastic games:
  - "Curse of dimensionality" in computing players' expectations over all possible future states of the game:
    - Suppose that each of N players can move to one of K states from one period to the next.
    - Need to sum up K^N terms.
  - Computational burden increases exponentially in the number of state variables $\rightarrow$ limited range and richness of applications.

- Continuous-time, finite-state stochastic games:
  - Computational advantages:
    - Avoids the curse of dimensionality (under widely used laws of motion).
    - Exploits precomputed addresses.
  - Conceptual differences:
    - Embeddability.
    - State changes.
    - Strategic interactions.
    - Calendar time and deterministic transitions.

Learning-by-Doing

- Besanko, D., Doraszelski, U., Kryukov, S. & Satterthwaite, M. (2010) "Learning-by-Doing, Organizational Forgetting, and Industry Dynamics."

- Borkovsky, R., Doraszelski, U., & Kryukov, Y. (2010) "A User's Guide to Solving Dynamic Stochastic Games Using the Homotopy Method."

- Question: Is organizational forgetting an antidote to market dominance?

- Incorporate organizational forgetting into the Cabral & Riordan (1994) model of learning-by-doing.

- Dynamic competition with learning-by-doing and organizational forgetting is akin to racing down an upward-moving escalator.

- Organizational forgetting makes bidirectional movements through the state space possible. Thus, it is a source of...
  - ...aggressive pricing behavior;
  - ...market dominance;
  - ...multiple equilibria.

- Learning-by-doing and organizational forgetting are distinct economic forces.

- Show that there are equilibria that the Pakes & McGuire (1994) algorithm cannot compute.

- Propose a homotopy algorithm to trace out the equilibrium correspondence.

Setup and Timing

- Discrete time, infinite horizon.
- Two firms with potentially different stocks of know-how
$$\omega = (\omega_1, \omega_2) \in \{1, \ldots, L\}^2 = \Omega.$$
- In each period, the timing is as follows:
  - Firms choose prices.
  - One buyer enters the market and makes at most one purchase.
  - Learning-by-doing and organizational forgetting occur and the firms' stocks of know-how change accordingly.
- Law of motion:
$$\omega_n' = \omega_n + q_n - f_n,$$
where
  - $q_n \in \{0, 1\}$ indicates whether firm n makes a sale with
$$\Pr(q_n = 1) = D_n(p_1, p_2) = \frac{\exp(v - p_n)}{1 + \sum_{k=1}^{2} \exp(v - p_k)};$$
  - $f_n \in \{0, 1\}$ represents organizational forgetting with
$$\Pr(f_n = 1) = \Delta(\omega_n) = 1 - (1 - \delta)^{\omega_n}.$$

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

$$V_n(\omega) = \max_{p_n} D_n(p_n, p_{-n}(\omega))(p_n - c(\omega_n)) + \beta \sum_{k=0}^{2} D_k(p_n, p_{-n}(\omega)) W_{nk}(\omega),$$

  where

  - $p_{-n}(\omega)$ is the price charged by the other firm;
  - the marginal cost of production is

$$c(\omega_n) = \begin{cases} \kappa \omega_n^{\eta} & \text{if} \quad 1 \leq \omega_n < l, \\ \kappa l^{\eta} & \text{if} \quad l \leq \omega_n \leq L, \end{cases}$$

    with $\eta = \log_2 \rho$ for a progress ratio of ρ;
  - $\beta \in (0, 1)$ is the discount factor;
  - $W_{nk}(\omega)$ is the expectation of firm n's value function conditional on buyer purchasing good $k \in \{0, 1, 2\}$ (good 0 is outside good).

Bellman Equation

- Continuation values:

$$W_{n0}(\omega) = \sum_{\omega_1'=1}^{L} \sum_{\omega_2'=1}^{L} V(\omega')\Pr(\omega_1'|\omega_1, q_1 = 0)\Pr(\omega_2'|\omega_2, q_2 = 0),$$

$$W_{n1}(\omega) = \sum_{\omega_1'=1}^{L} \sum_{\omega_2'=1}^{L} V(\omega')\Pr(\omega_1'|\omega_1, q_1 = 1)\Pr(\omega_2'|\omega_2, q_2 = 0),$$

$$W_{n2}(\omega) = \sum_{\omega_1'=1}^{L} \sum_{\omega_2'=1}^{L} V(\omega')\Pr(\omega_1'|\omega_1, q_1 = 0)\Pr(\omega_2'|\omega_2, q_2 = 1),$$

where

$$\Pr(\omega_n'|\omega_n, q_n) = \begin{cases} 1 - \Delta(\omega_n) & \text{if} \quad \omega_n' = \omega_n + q_n, \\ \Delta(\omega_n) & \text{if} \quad \omega_n' = \omega_n + q_n - 1, \end{cases}$$

and $\Pr(L|L, q_n = 1) = 1$ and $\Pr(1|1, q_n = 0) = 1$.

- $p_n(\omega)$ is unique solution to FOC:

$$0 = 1 - (1 - D_n(p_n, p_{-n}(\omega)))\,(p_n - c(\omega_n)) - \beta W_{nn}(\omega)$$
$$+\beta \sum_{k=0}^{2} D_k(p_n, p_{-n}(\omega)) W_{nk}(\omega).$$

Equilibrium

- Primitives are 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 $p_1(\omega_1, \omega_2) = p(\omega_1, \omega_2)$ and $p_2(\omega_1, \omega_2) = p(\omega_2, \omega_1)$.

- Bellman equation and FOC for state ω are

$$V(\omega) = D_1(\omega)\left(p(\omega) - c(\omega_1)\right) + \beta \sum_{k=0}^{2} D_k(\omega) W_k(\omega),$$

$$0 = 1 - \left(1 - D_1(\omega)\right)\left(p(\omega) - c(\omega_1)\right) - \beta W_1(\omega) + \beta \sum_{k=0}^{2} D_k(\omega) W_k(\omega),$$

  where $D_k(\omega) = D_k(p(\omega), p(\omega_2, \omega_1))$, $k \in \{0, 1, 2\}$.

- The system of $2L^2$ nonlinear equations given by the collection of the above equations for each state $\omega \in \{1, \ldots, L\}^2$ defines a symmetric equilibrium.

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

Computation: Homotopy Algorithm

- Write the system of $2L^2$ nonlinear equations (Bellman equations and FOCs) as
$$\mathbf{F}(\mathbf{x}, \delta) = 0,$$
where
$$\mathbf{x} = (V(1,1), \ldots, V(L,L), p(1,1), \ldots, p(L,L)).$$

- The object of interest is the equilibrium correspondence
$$\mathbf{F}^{-1} = \{(\mathbf{x}, \delta) | \mathbf{F}(\mathbf{x}, \delta) = 0\}.$$

- The homotopy algorithm follows a path from the unique equilibrium at $\delta = 0$ to the unique equilibrium at $\delta = 1$.

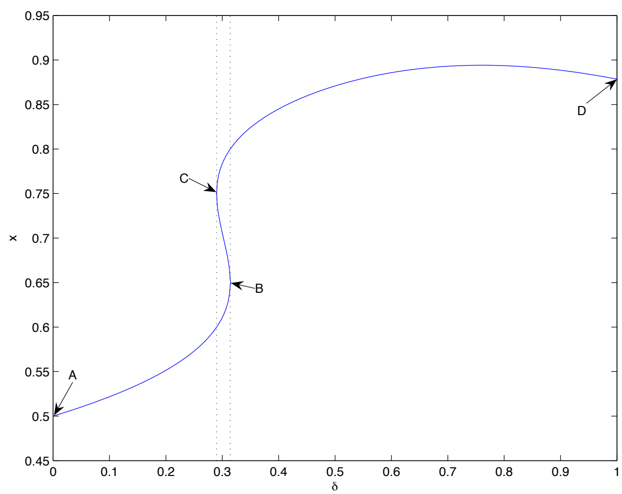

Example.

Computation: Homotopy Algorithm

- Define a parametric path to be a set of functions $(\mathbf{x}(s), \delta(s))$ such that $(\mathbf{x}(s), \delta(s)) \in \mathbf{F}^{-1}$.

- The conditions that are required to remain "on path" are found by differentiating

$$\mathbf{F}(\mathbf{x}(s), \delta(s)) = 0$$

with respect to s:

$$\sum_{i=1}^{2L^2} \frac{\partial \mathbf{F}(\mathbf{x}(s), \delta(s))}{\partial x_i} x_i'(s) + \frac{\partial \mathbf{F}(\mathbf{x}(s), \delta(s))}{\partial \delta} \delta'(s) = 0.$$

- While there are many solutions, all of them describe the same path.

- One solution obeys the so-called basic differential equations (BDE)

$$y_i'(s) = (-1)^{i+1} \det \left(\left(\frac{\partial \mathbf{F}(\mathbf{y}(s))}{\partial \mathbf{y}} \right)_{-i} \right), \quad i = 1, \ldots, 2L^2 + 1, \tag{4}$$

where $\mathbf{y}(s) = (\mathbf{x}(s), \delta(s))$ and the notation $(\cdot)_{-i}$ is used to indicate that the ith column is removed from the $(2L^2 \times 2L^2 + 1)$ Jacobian $\frac{\partial \mathbf{F}(\mathbf{y}(s))}{\partial \mathbf{y}}$.

- The BDE reduce the task of tracing out the equilibrium correspondence to solving a system of differential equations.

Computation: Pakes & McGuire (1994) Algorithm

Executes the iteration

$$\mathbf{x}^{l+1} = \mathbf{G}(\mathbf{x}^l), \quad l = 0, 1, 2, \ldots,$$

where, for each state $\omega \in \{1, \ldots, L\}^2$, old guesses for the value and policy of firm 1 are mapped into new guesses as follows:

$$\begin{aligned}
p^{l+1}(\omega) &= \arg\max_{p_1} D_1(p_1, p^l(\omega_2, \omega_1))(p_1 - c(\omega_1)) \\
&\quad + \beta \sum_{k=0}^{2} D_k(p_1, p^l(\omega_2, \omega_1)) W_k^l(\omega), \\
V^{l+1}(\omega) &= D_1(p^{l+1}(\omega), p^l(\omega_2, \omega_1)) \left(p^{l+1}(\omega) - c(\omega_1)\right) \\
&\quad + \beta \sum_{k=0}^{2} D_k(p^{l+1}(\omega), p^l(\omega_2, \omega_1)) W_k^l(\omega).
\end{aligned}$$

Computation: Pakes & McGuire (1994) Algorithm

- “Inbetween” two equilibria that can be computed using the Pakes & McGuire (1994) algorithm, there is one equilibrium that cannot:

  Proposition 1 *If $\delta'(s) \leq 0$, then $\varrho\left(\left.\frac{\partial \mathbf{G}(\mathbf{x}(s))}{\partial \mathbf{x}}\right|_{\delta(s)}\right) \geq 1$.*

- Let I denote the $(2L^2 \times 2L^2)$ identity matrix. Then

$$\left.\frac{\partial \mathbf{G}(\mathbf{x}(s))}{\partial \mathbf{x}}\right|_{\delta(s)} = \frac{\partial \mathbf{F}(\mathbf{x}(s), \delta(s))}{\partial \mathbf{x}} + I. \tag{5}$$

- The BDE (4) imply

$$\delta'(s) = \det\left(\frac{\partial \mathbf{F}(\mathbf{x}(s), \delta(s))}{\partial \mathbf{x}}\right).$$

- Since the determinant of $\frac{\partial \mathbf{F}(\mathbf{x}(s),\delta(s))}{\partial \mathbf{x}}$ is the product of $2L^2$ eigenvalues, if $\delta'(s) \leq 0$, then there exists at least one real nonnegative eigenvalue.

- Let A be an arbitrary matrix and $\varsigma(A)$ its spectrum. Then $\varsigma(A + I) = \varsigma(A) + 1$.

- It follows from equation (5) that $\left.\frac{\partial \mathbf{G}(\mathbf{x}(s))}{\partial \mathbf{x}}\right|_{\delta(s)}$ has at least one real eigenvalue equal to or bigger than unity.

Equilibrium Correspondence: Multiple Equilibria

Proposition 2 *If organizational forgetting is either absent ($\delta = 0$) or certain ($\delta = 1$), then there is a unique equilibrium.*

Result 1 *If organizational forgetting is neither absent ($\delta = 0$) nor certain ($\delta = 1$), then there may be multiple equilibria.*

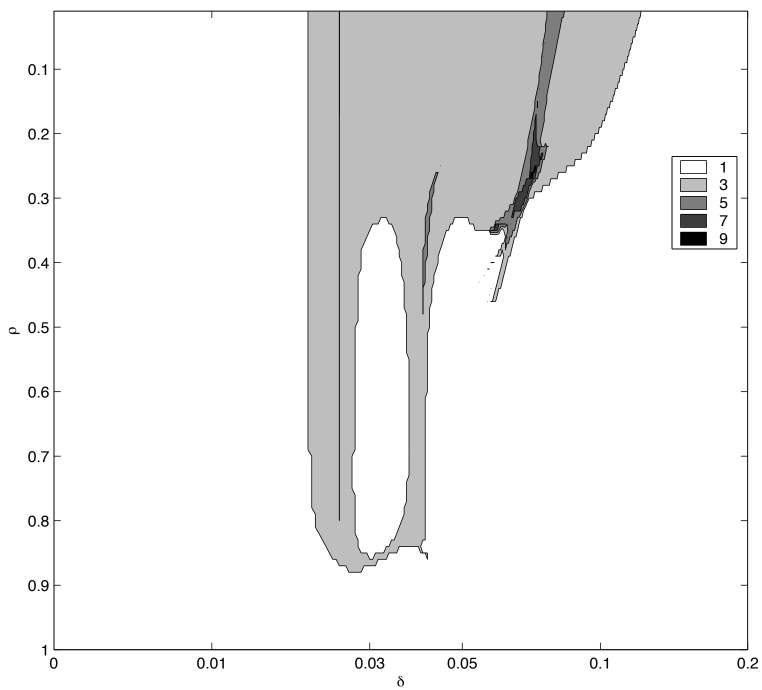

Number of equilibria.

Equilibrium Correspondence: Paths and Loops

Result 2 *The equilibrium correspondence* $\mathbf{F}^{-1}$ *contains a unique path that connects the equilibrium at* $\delta = 0$ *with the equilibrium at* $\delta = 1$*. In addition,* $\mathbf{F}^{-1}$ *may contain (one or more) loops that are disjoint from this "main path" and from each other.*

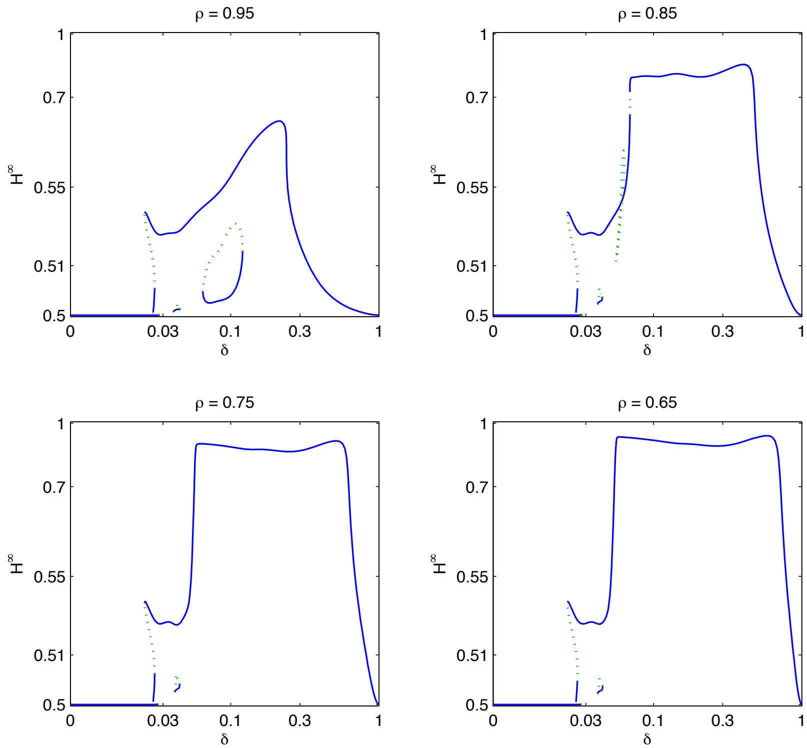

Limiting expected Herfindahl index H^{∞}. Equilibria with $\varrho\left(\frac{\partial \mathbf{G}(\mathbf{x}(s))}{\partial \mathbf{x}}\Big|_{\delta(s)}\right) < 1$ (solid line) and equilibria with $\varrho\left(\frac{\partial \mathbf{G}(\mathbf{x}(s))}{\partial \mathbf{x}}\Big|_{\delta(s)}\right) \geq 1$ (dotted line).

Equilibrium Correspondence: Categories of Equilibria

Categories of equilibria:

- flat without well;
- flat with well;
- trenchy;
- extra-trenchy.

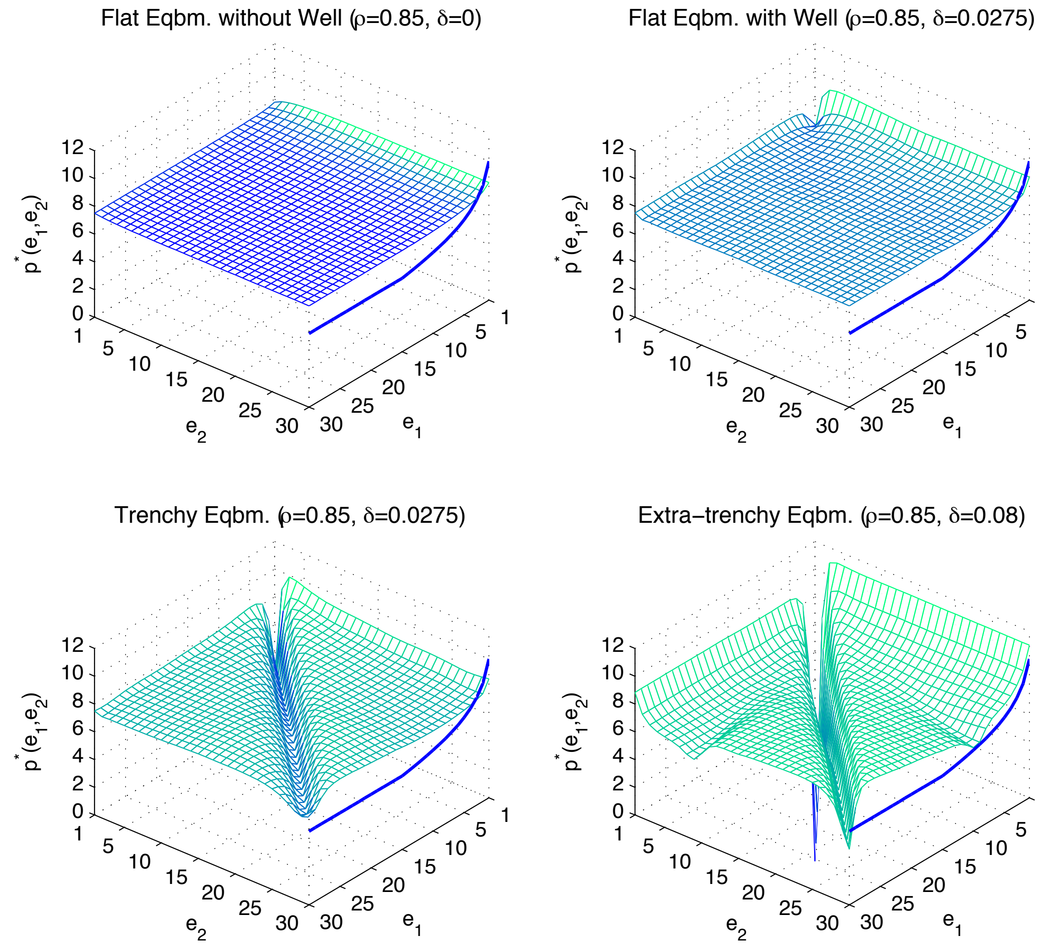

Policy function $p^*(\omega_1, \omega_2)$. Marginal cost $c(\omega_1)$ (solid line in $\omega_2 = 30$-plane).

Industry Dynamics: Flat Equilibrium with Well ($\rho = 0.85$, $\delta = 0.03$)

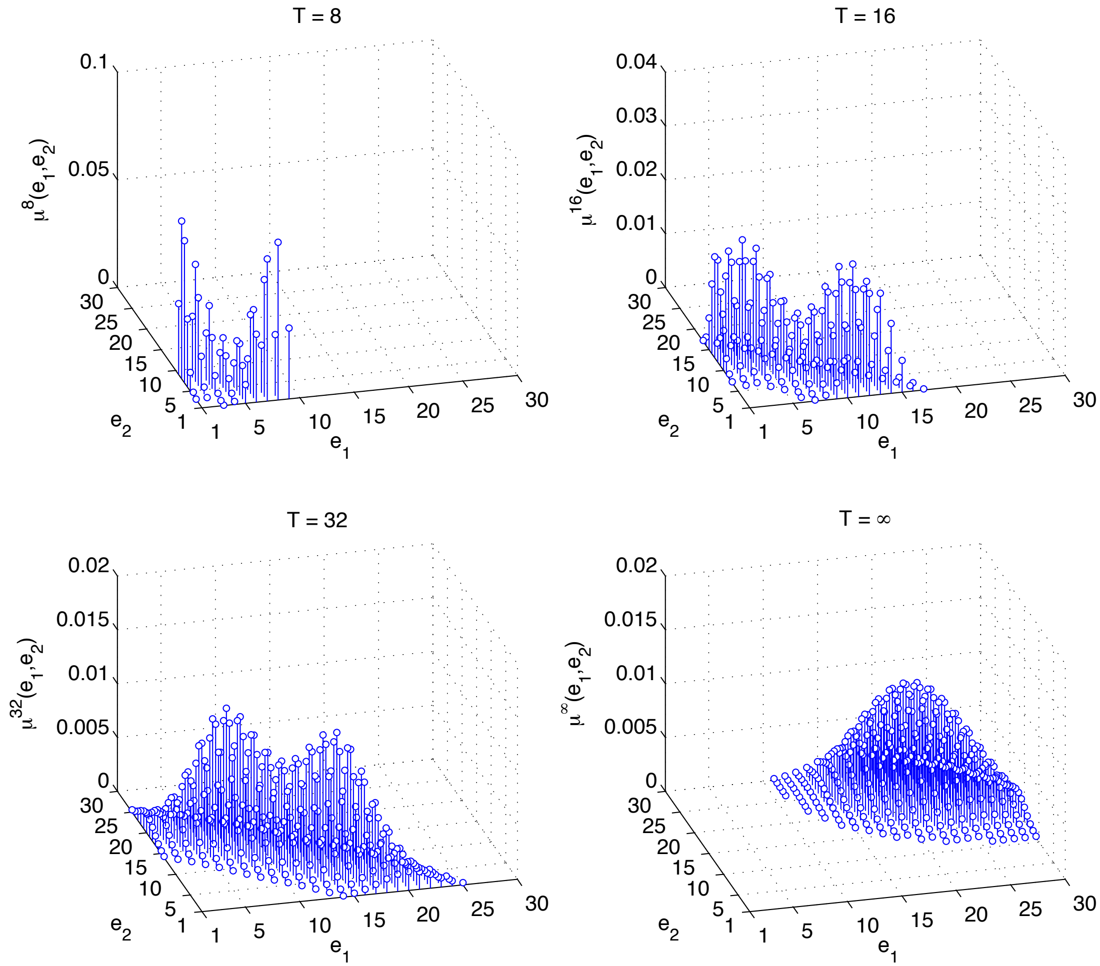

Transient distribution over states in periods 8, 16, and 32 given initial state $(1,1)$ and limiting distribution.

Industry Dynamics: Trenchy Equilibrium ($\rho = 0.85$, $\delta = 0.03$)

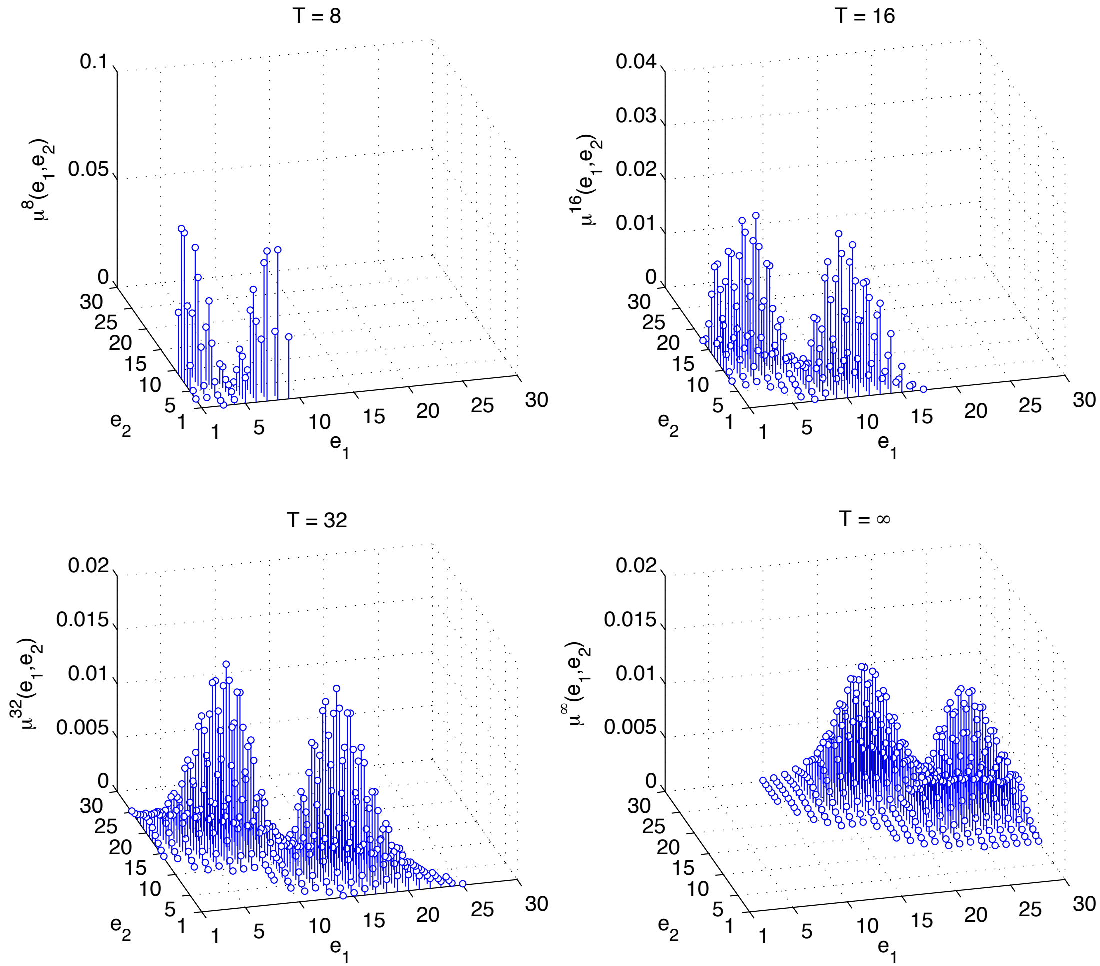

Transient distribution over states in periods 8, 16, and 32 given initial state $(1, 1)$ and limiting distribution.

Organizational Forgetting and Multiple Equilibria

- What gives rise to multiple equilibria ranging from "peaceful coexistence" to "trench warfare"?

- Holding the value of continued play fixed, the strategic situation in state ω is akin to a static game.

  Proposition 3 *Statewise uniqueness holds provided the outside good is sufficiently unattractive (υ large).*

- Multiple equilibria must arise from firms' expectations regarding the value of continued play.

  Taking the value of continued play as given, the reaction functions intersect once, but there is more than one value of continued play that is consistent with rational expectations.

- Multiplicity is rooted in the dynamics of the model.

Organizational Forgetting and Multiple Equilibria

- When do multiple equilibria arise?

- In expectation, the "inflow" of know-how into the industry is almost one unit per period, the "outflow" in state ω is $\Delta(\omega_1) + \Delta(\omega_2)$.

- Consider state (ω, ω), where $\omega \geq l$.
  - If $1 \ll 2\Delta(\omega)$, then it is virtually impossible that both firms reach the bottom of their learning curves $\rightarrow$ trench warfare.
  - If $1 \gg 2\Delta(\omega)$, then it is virtually inevitable that both firms reach the bottom of their learning curves $\rightarrow$ peaceful coexistence.
  - If $1 \approx 2\Delta(\omega)$, then primitives do not suffice to tie down the equilibrium $\rightarrow$ multiple equilibria.

    Back-of-the-envelope calculation ($l = 15$ and $L = 30$):

$$1 = 2\Delta(15) \Rightarrow \delta = 0.05 \text{ and } 1 = 2\Delta(30) \Rightarrow \delta = 0.02.$$

- Stagewise uniqueness and unidirectional movements through the state space $\rightarrow$ unique equilibrium.

  Organizational forgetting makes bidirectional movements possible.