Dynamic Games: Numerical Methods and Applications

Day 2b

Prof. Ulrich Doraszelski, University of Pennsylvania

Agenda

- Computational burden.
- Oblivious equilibrium.
- Continuous-time stochastic games.
- Discrete-time stochastic games with alternating moves.
- Open questions.

Sources of Computational Burden

- State space:
  - Suppose that each of N players can be at one of L states.
  - State space has L^N elements.
  - Symmetry reduces exponential to polynomial growth.

- Successor states:
  - Suppose that each of N players can move to one of K states from one period to the next.
  - Expectation over successor states involves K^N terms.

Alleviating the Computational Burden

System of equations:

- Ferris, M., Judd, K. & Schmedders, K. (2007) "Solving Dynamic Games with Newton's Method."

Ergodic set:

- Pakes, A. & McGuire, P. (2001) "Stochastic Algorithms, Symmetric Markov Perfect Equilibrium, and the 'Curse' of Dimensionality."

- Judd, K., Maliar, L. & Maliar, S. (2012) "Merging Simulation and Projection Approaches to Solve High-Dimensional Problems."

State aggregation and interpolation methods:

- 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."

Alleviating the Computational Burden

Oblivious equilibrium and its extensions:

- 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."
- Weintraub, G., Benkard, L. Jeziorski, P. & Van Roy, B. (2008) "Nonstationary Oblivious Equilibrium."
- Benkard, L., Jeziorski, P. & Weintraub, G., (2012) "Oblivious Equilibrium for Concentrated Industries."
- Ifrach, B. and Weintraub, G. (2012) "A Framework for Dynamic Oligopoly in Concentrated Industries."

Alleviating the Computational Burden

Continuous-time stochastic games:

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

- Arcidiacono, P. Bayer, P. Blevins, J. & Ellickson (2013) "Estimation of Dynamic Discrete Choice Models in Continuous Time with an Application to Retail Competition."

Discrete-time stochastic games with alternating moves:

- Doraszelski, U. & Judd, K. (2007) "Dynamic Stochastic Games with Sequential State-to-State Transitions."

Oblivious Equilibrium

- 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."

- Markovian strategy depends on focal firm's state and rivals' states or, with symmetry, distribution over rivals' states.

- Oblivious strategy depends on focal firm's state and long-run expected industry state.
  - Focal firm presumes that its rivals' state is long-run expected industry state forever $\rightarrow$ limits aggregate shocks and transitional dynamics.
  - Focal firm ignores its impact on long-run expected industry state $\rightarrow$ returns to literature on long-run industry equilibrium (Jovanovic 1982, Hopenhayn 1992, Melitz 2003).

    Additional reading:
    * Weintraub, G., Benkard, L. & Van Roy, B. (2011) "Industry Dynamics: Foundations for Models with an Infinite Number of Firms."

Oblivious Equilibrium: Main Idea

- Make initial guess for oblivious strategy.

- Compute long-run expected industry state induced by this strategy.

- Solve focal firm's dynamic programming problem: Find oblivious strategy that maximizes focal firm's expected NPV given the long-run expected industry state.

- Repeat until just-computed oblivious strategy is close to initially-guessed oblivious strategy.

Oblivious Equilibrium as Approximation to Markov-Perfect Equilibrium

- Oblivious equilibrium approximates Markov-perfect equilibrium if

  1. the average number of firms grows to infinity as market size grows $\rightarrow$ industry state is almost constant;

  2. firm size distribution is "light-tailed" $\rightarrow$ no reason to keep track of dominant firm.

- Error bounds: How much can focal firm gain, on average, by deviating to Markovian strategy if it's rivals adhere to oblivious strategy?

Moment-Based Markov Equilibrium

- Ifrach, B. and Weintraub, G. (2012) "A Framework for Dynamic Oligopoly in Concentrated Industries."

- Few dominant firms, many fringe firms.

- Moment-based strategy depends on focal firm's state, dominant firms' states, and moments of distribution over fringe firms' states (state aggregation).

- Plausible description of behavior if there are many fringe firms.

- Conceptual issues:

  1. Per-period profit may depend on distribution over fringe firms' states, and not just moments.

  2. Moments are not Markovian even if states are $\rightarrow$ moments are not sufficient statistics for industry evolution.

     Example: Moment is average. Same average if one fringe firm is in state 2 and one in state 6 and if both are in state 4, but potentially very different behaviors.

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 and limitations:
    * Embeddability.
    * State changes.
    * Strategic interactions.
    * Calendar time and deterministic transitions.

Continuous-Time Model

- Set of possible states Ω is finite. State at time t is ω_t. In applications ω_t is a vector partitioned into $(\omega_t^1, \ldots, \omega_t^N)$.

- N players. Player i's action in period t is $x_t^i \in \mathbb{X}^i(\omega_t)$. Set of feasible actions $\mathbb{X}^i(\omega_t)$ is arbitrary.

  Notation: $x_t = \left(x_t^1, \ldots, x_t^N\right)$ and $x_t^{-i} = \left(x_t^1, \ldots, x_t^{i-1}, x_t^{i+1}, \ldots, x_t^N\right)$.

- Law of motion: State follows a controlled continuous-time, finite-state Markov process.
  - Time path is a piecewise-constant, right-continuous function of time.
  - Hazard rate of a jump occurring is $\phi(x_t, \omega_t)$.
  - Conditional on a jump occurring, transition probability is $f\left(\omega' | \omega_{t^-}, x_{t^-}\right)$, where $\omega_{t^-} = \lim_{s \to t^-} \omega_s$ and $x_{t^-} = \lim_{s \to t^-} x_s$.

Continuous-Time Model (cont'd)

- Law of motion (cont'd): Over a short interval of time of length $\Delta > 0$ the law of motion is

$$\begin{aligned} \Pr\left(\omega_{t+\Delta} \neq \omega_t | \omega_t, x_t\right) &= \phi\left(x_t, \omega_t\right)\Delta + O\left(\Delta^2\right), \\ \Pr\left(\omega_{t+\Delta} = \omega' | \omega_t, x_t, \omega_{t+\Delta} \neq \omega_t\right) &= f\left(\omega' | \omega_t, x_t\right) + O(\Delta). \end{aligned}$$

  Special case of independent transitions:

$$\begin{aligned} \Pr^i\left(\omega^i_{t+\Delta} \neq \omega^i_t | \omega^i_t, x^i_t\right) &= \phi^i\left(x^i_t, \omega^i_t\right)\Delta + O\left(\Delta^2\right), \\ \Pr^i\left(\omega^i_{t+\Delta} = \left(\omega'\right)^i | \omega^i_t, x^i_t, \omega^i_{t+\Delta} \neq \omega^i_t\right) &= f^i\left(\left(\omega'\right)^i | \omega^i_t, x^i_t\right) + O(\Delta), \end{aligned}$$

  and $\phi\left(x_t, \omega_t\right) = \sum_{i=1}^N \phi^i\left(x^i_t, \omega^i_t\right)$.

  During a short interval of time, there will be (with probability infinitesimally close to one) at most one jump.

- Payoffs: Player i receives a payoff flow of $\pi^i\left(x_t, \omega_t\right)$ when players' actions are x_t and the state is ω_t.

  Objective is to maximize the expected NPV of future cash flows

$$\mathrm{E}\left\{\int_0^\infty e^{-\rho t}\pi^i\left(x_t, \omega_t\right)dt\right\},$$

  where $\rho > 0$ is the discount rate.

Continuous-Time Bellman Equation

- Over a short interval of time of length $\Delta > 0$ player i solves the dynamic programming problem given by

$$
\begin{aligned}
V^i(\omega) \quad = \quad & \max_{x^i} \pi^i\left(x^i, X^{-i}(\omega), \omega\right)\Delta \\
& + (1-\rho\Delta)\Big\{ \left(1 - \phi\left(x^i, X^{-i}(\omega), \omega\right)\Delta - O\left(\Delta^2\right)\right) V^i(\omega) \\
& + \left(\phi\left(x^i, X^{-i}(\omega), \omega\right)\Delta + O\left(\Delta^2\right)\right)\left(\mathsf{E}_{\omega'}\left\{V^i\left(\omega'\right)|\omega, x^i, X^{-i}(\omega)\right\} + O(\Delta)\right)\Big\}.
\end{aligned}
$$

- As $\Delta \to 0$, the Bellman equation of player i becomes

$$
\begin{aligned}
\rho V^i(\omega) \quad = \quad & \max_{x^i} \pi^i\left(x^i, X^{-i}(\omega), \omega\right) - \phi\left(x^i, X^{-i}(\omega), \omega\right) V^i(\omega) \\
& + \phi\left(x^i, X^{-i}(\omega), \omega\right) \mathsf{E}_{\omega'}\left\{V^i\left(\omega'\right)|\omega, x^i, X^{-i}(\omega)\right\}. \qquad (1)
\end{aligned}
$$

  Special case of independent transitions:

$$
\begin{aligned}
\rho V^i(\omega) \quad = \quad & \max_{x^i} \pi^i\left(x^i, X^{-i}(\omega), \omega\right) - \phi\left(x^i, X^{-i}(\omega), \omega\right) V^i(\omega) \\
& + \phi^i\left(x^i, \omega^i\right) \mathsf{E}_{(\omega')^i}\left\{V^i\left(\left(\omega'\right)^i, \omega^{-i}\right)|\omega^i, x^i\right\} \\
& + \sum_{j \neq i} \phi^j\left(X^j(\omega), \omega^j\right) \mathsf{E}_{(\omega')^j}\left\{V^i\left(\left(\omega'\right)^j, \omega^{-j}\right)|\omega^j, X^j(\omega))\right\}. \quad (2)
\end{aligned}
$$

- Player i's strategy is found by carrying out the maximization on the RHS of equation (1) or (2).

Avoiding the Curse of Dimensionality

- Consider the special case of independent transitions.

- Continuous-time model requires computing N one-dimensional expectations instead of one N-dimensional expectation:

$$\mathsf{E}_{(\omega')^j}\left\{V^i\left(\left(\omega'\right)^j,\omega^{-j}\right)|\omega^j,X^j(\omega)\right\}=\sum_{(\omega')^j\in\{\omega^j-1,\omega^j+1\}}V^i\left(\left(\omega'\right)^j,\omega^{-j}\right)f^j\left(\left(\omega'\right)^j|\omega^j,X^j(\omega)\right).$$

  Expectation involves $(K-1)N$ terms.

- Discrete-time model requires computing the expectation

$$\mathsf{E}_{\omega'}\left\{V^i\left(\omega'\right)|\omega,X(\omega)\right\}=\sum_{\left\{\omega':(\omega')^i\in\{\omega^i-1,\omega^i,\omega^i+1\},i=1,\ldots,N\right\}}V^i\left(\omega'\right)\prod_{i=1}^N\mathsf{Pr}^i\left(\left(\omega'\right)^i|\omega^i,X^i(\omega)\right).$$

  If each of N players can move to one of K states, then expectation over successor states involves K^N terms.

- What matters is the total number of coordinates of the state vector:
  - Multiple states per player.
  - Common states and common shocks.

Continuous-Time Algorithm

- Block Gauss-Seidel scheme: Algorithm is iterative and proceeds through the state space Ω in some prespecified order.

- Given old guesses $V^i(\omega)$ and $X^i(\omega)$ it computes new guesses $\widehat{V}^i(\omega)$ and $\widehat{X}^i(\omega)$ for each player $i = 1, \ldots, N$ as follows:

$$\begin{aligned} \widehat{X}^i(\omega) \quad &\leftarrow \quad \arg\max_{x^i} \pi^i\left(x^i, X^{-i}(\omega), \omega\right) - \phi\left(x^i, X^{-i}(\omega), \omega\right) V^i(\omega) \\ &\quad +\phi\left(x^i, X^{-i}(\omega), \omega\right) \mathsf{E}_{\omega'}\left\{V^i\left(\omega'\right) | \omega, x^i, X^{-i}(\omega)\right\}, \\ \widehat{V}^i(\omega) \quad &\leftarrow \quad \frac{1}{\rho + \phi\left(\widehat{X}^i(\omega), X^{-i}(\omega), \omega\right)} \pi^i\left(\widehat{X}^i(\omega), X^{-i}(\omega), \omega\right) \\ &\quad + \frac{\phi\left(\widehat{X}^i(\omega), X^{-i}(\omega), \omega\right)}{\rho + \phi\left(\widehat{X}^i(\omega), X^{-i}(\omega), \omega\right)} \mathsf{E}_{\omega'}\left\{V^i\left(\omega'\right) | \omega, \widehat{X}^i(\omega), X^{-i}(\omega)\right\}. \end{aligned}$$

- Dividing through by $\rho + \phi(\widehat{X}^i(\omega), X^{-i}(\omega), \omega)$ ensures that equation is contractive for a given player (holding fixed the policies of all players) since

$$\frac{\phi(\widehat{X}^i(\omega), X^{-i}(\omega), \omega)}{\rho + \phi(\widehat{X}^i(\omega), X^{-i}(\omega), \omega)} < 1.$$

Precomputed Addresses, Symmetry, and Anonymity

- Store the value and policy functions in table $\mathbb{M}$. Each row corresponds to a state $\omega \in \Omega$ and contains the vector
$$\left(V^1(\omega), \ldots, V^N(\omega), X^1(\omega), \ldots, X^N(\omega)\right).$$

- Expectation $\mathsf{E}_{\omega'}\left\{V^i\left(\omega'\right)|\omega, X(\omega)\right\}$ is really
$$\sum_{\{\omega' : \Pr(\omega'|\omega, \mathbb{M}[R(\omega),(N+1,\ldots,2N)])>0\}} \mathbb{M}\left[R\left(\omega'\right), C\left(\omega', i\right)\right] \Pr\left(\omega'|\omega, \mathbb{M}\left[R\left(\omega\right), \left(N+1, \ldots, 2N\right)\right]\right),$$
where $C\left(\omega', i\right)$ is the column in row $R\left(\omega'\right)$ that contains the value for player i in state ω', etc.

- Two kinds of costs:
  - Summation over all states ω' such that $\Pr\left(\omega'|\omega, X(\omega)\right) > 0$.
  - Computation of the address, $R(\omega')$ and $C(\omega', i)$, of the value of player i at each of them.

- Idea: Precompute and store addresses of successor states.

- Feasible if number of successor states is small $\rightarrow$ continuous time.

- Useful if $R(\omega)$ and $C(\omega, i)$ are hard to evaluate $\rightarrow$ symmetry and anonymity.

Precomputed Addresses, Symmetry, and Anonymity (cont'd)

- Set of possible states of player i is $\{1,\ldots,M\} \to \Omega = \{1,\ldots,M\}^N$.

- Invoke symmetry and anonymity to slow down the growth of the state space in N and M.

  – Symmetry allows us to focus on the problem of player 1. Formally,

  $$V^i(\omega) = V^1(\omega^i, \omega^2, \ldots, \omega^{i-1}, \omega^1, \omega^{i+1}, \ldots).$$

  – Anonymity (also called exchangeability) says that player 1 does not care about the identity of its competitors. Formally,

  $$V^1(\omega) = V^1(\omega^1, \ldots, \omega^{j-1}, \omega^k, \omega^{j+1}, \ldots, \omega^{k-1}, \omega^j, \omega^{k+1}, \ldots)$$

  for all $2 \leq j < k$.

- In practice symmetry and anonymity are imposed by focusing on

  $$\bar{\Omega} = \{(\omega_1, \omega_2, \ldots, \omega_N) \in \Omega : \omega_1 \leq \omega_2 \leq \ldots \leq \omega_N\}.$$

  – Reduces the number of states to be examined from $|\Omega| = M^N$ to $|\bar{\Omega}| = \frac{(N+M-1)!}{N!(M-1)!}$.

  – Makes $R(\omega)$ and $C(\omega, i)$ much harder to compute (see Pakes, Gowrisankaran & McGuire (1993) and Gowrisankaran (1999) for details).

Example: The Pakes & McGuire (1994) Quality Ladder Model

- Firm i produces a product of quality $\omega^i \in \{1, \ldots, M\}$.
- Heterogeneous consumers with unit demands. Firm i's demand is
$$q^i(p^1, \ldots, p^N; \omega) = m \frac{\exp\left(g(\omega^i) - p^i\right)}{1 + \sum_{j=1}^N \exp\left(g(\omega^j) - p^j\right)},$$
where $m > 0$ is market size.
- Firm i sets its price p^i to maximize profits:
$$\max_{p^i \geq 0} q^i(p^1, \ldots, p^N; \omega)\left(p^i - c\right),$$
where $c \geq 0$ is marginal cost of production.
- Payoff function:
$$\pi^i(x, \omega) \equiv q^i(p^1(\omega), \ldots, p^N(\omega); \omega)(p^i(\omega) - c) - x^i,$$
where $x^i \geq 0$ is firm i's investment in quality improvements.
- Law of motion:
  - Successful investment has probability $\frac{\alpha x^i}{1+\alpha x^i}$.
  - Depreciation shock has probability δ.
- Parameterization: See Pakes & McGuire (1994).

Example: Reformulating the Model in Continuous Time

- Payoff function: Same as in discrete-time model.

- Law of motion:

  - Success hazard is $\frac{\alpha x^i}{1+\alpha x^i}$.

  - Depreciation hazard is δ.

  Hazard rate of a jump in firm i's state occurring is

$$\phi^i(x^i, \omega^i) = \frac{\alpha x^i}{1 + \alpha x^i} + \delta.$$

  Conditional on a jump occurring, transition probability is

$$f^i\left(\left(\omega'\right)^i | \omega^i, x^i\right) = \begin{cases} \frac{\alpha x^i}{(1+\alpha x^i)\phi^i(x^i,\omega^i)}, & (\omega')^i = \omega^i + 1, \\ \frac{\delta}{\phi^i(x^i,\omega^i)}, & (\omega')^i = \omega^i - 1. \end{cases}$$

- Parameterization: If Δ is unit of time in discrete-time model, then $\beta = e^{-\rho\Delta} \Leftrightarrow \rho = -\frac{\ln\beta}{\Delta}$. Take $\Delta = 1$ to obtain $\rho = -\ln\beta$.

Computational Advantages

- Time per iteration.
  - Avoiding the curse of dimensionality.
  - Precomputed addresses.
- Number of iterations.
  - Iteration penalty.
- Time to convergence.

Computational Advantages: Time per Iteration

#firms	#states	#unknowns	discrete time secs.	discrete time perct.	continuous time without precomputed addresses secs.	continuous time without precomputed addresses perct.	continuous time with precomputed addresses secs.	continuous time with precomputed addresses perct.	ratio: discrete to continuous time without precomputed addresses	ratio: continuous time without to with precomputed addresses	ratio: discrete to continuous time with precomputed addresses
2	45	180	9.78(-7)	52%	6.89(-7)	42%	5.67(-7)	33%	1.42	1.22	1.73
3	165	990	1.45(-6)	74%	6.36(-7)	44%	5.05(-7)	38%	2.29	1.26	2.88
4	495	3960	2.90(-6)	88%	6.36(-7)	48%	4.75(-7)	43%	4.55	1.34	6.10
5	1287	12870	6.94(-6)	96%	6.42(-7)	53%	4.77(-7)	46%	10.81	1.35	14.57
6	3003	36036	1.81(-5)	98%	6.88(-7)	55%	4.88(-7)	45%	26.34	1.41	37.12
7	6435	90090	5.02(-5)	100%	7.33(-7)	53%	5.11(-7)	48%	68.48	1.43	98.26
8	12870	205920	1.31(-4)	100%	7.77(-7)	55%	5.24(-7)	50%	168.33	1.48	249.38
9	24310	437580	3.82(-4)	100%	7.77(-7)	62%	5.39(-7)	53%	492.16	1.44	709.04
10	43758	875160	1.07(-3)	100%	8.34(-7)	64%	5.94(-7)	44%	1282.19	1.40	1800.00
11	75582	1662804	2.99(-3)	100%	8.42(-7)	67%	5.77(-7)	56%	3557.14	1.46	5187.50
12	125970	3023280	8.20(-3)	100%	8.60(-7)	68%	5.95(-7)	60%	9533.08	1.44	13770.00
13	203490	5290740	2.42(-2)	100%	9.22(-7)	69%	6.20(-7)	61%	26235.65	1.49	39033.56
14	319770	8953560	6.76(-2)	100%	9.53(-7)	72%	6.55(-7)	59%	70946.70	1.45	103195.27

Time per iteration per state per firm and percentage of time spent on computing the expectation. (k) is shorthand for $\times 10^k$. Quality ladder model with $M = 9$ quality levels per firm and a discount factor of 0.925.

Computational Advantages: Number of Iterations

- Continuous-time "contraction" factor:

$$\eta(X(\omega),\omega) = \frac{\phi\left(X(\omega),\omega\right)}{\rho + \phi\left(X(\omega),\omega\right)}.$$

  Varies with players' policies from state to state.

  Interpretation: $\eta(X(\omega),\omega)$ is the expected NPV of a dollar delivered at the next time the state changes. If $\rho \ll \phi(X(\omega),\omega) = 1$, then

$$\eta(X(\omega),\omega) = \frac{1}{\rho + 1} \approx 1 - \rho = 1 + \ln\beta \approx \beta.$$

- If ρ is small or if $\phi\left(X\left(\omega\right),\omega\right)$ is large (in particular, if N is large), then $\eta(X(\omega),\omega)$ is large and the contraction aspect is weak.

Computational Advantages: Time to Convergence

#firms	discrete time (mins.)	continuous time (mins.)	time per iteration	ratio number of iterations	time to convergence
2	1.80(-4)	1.12(-4)	1.73	0.93	1.61
3	1.42(-3)	8.83(-4)	2.88	0.56	1.60
4	1.13(-2)	4.43(-3)	6.10	0.42	2.54
5	8.78(-2)	1.70(-2)	14.57	0.36	5.18
6	6.42(-1)	5.34(-2)	37.12	0.32	12.03
7	4.44(0)	1.47(-1)	98.26	0.31	30.19
8	2.67(1)	3.56(-1)	249.38	0.30	74.94
9	1.66(2)	7.95(-1)	709.04	0.29	208.85
10	*9.28(2)*	1.77(0)	1800.00	*0.29*	*523.72*
11	*4.94(3)*	3.30(0)	5187.50	*0.29*	*1498.33*
12	*2.46(4)*	6.18(0)	13770.00	*0.29*	*3977.26*
13	*1.27(5)*	1.13(1)	39033.56	*0.29*	*11246.96*
14	*6.00(5)*	2.02(1)	103195.27	*0.29*	*29734.23*

Time to convergence. (k) is shorthand for $\times 10^k$. Convergence criterion is "distance to truth$< 10^{-4}$." Entries in italics are based on an estimated 119 iterations to convergence in discrete time. Quality ladder model with $M = 9$ quality levels per firm and a discount factor of 0.925.

Computational Advantages: Summary

- Time per iteration.
  - Gain from breaking the curse of dimensionality: up to five orders of magnitude.
  - Gain from precomputed addresses: 20% to 50%.
- Number of iterations.
  - Loss from iteration penalty: less than one order of magnitude.
- Time to convergence.
  - Huge gain!!!

Discrete-Time Stochastic Games with Alternating Moves and Finite States

- Doraszelski, U. & Judd, K. (2007) "Dynamic Stochastic Games with Sequential State-to-State Transitions."

- Each period one player is picked at random to choose an action. The state of the player with the move then changes.

- Random-leadership Stackelberg game.

- Can get computational advantages similar to continuous time within the more familiar discrete-time framework.

Alternating-Moves Model

- Payoffs: Player i receives a per-period payoff of $\pi^{i,j}(x_t, \omega_t)$ when player j has the move and players' actions are x_t and the state is ω_t.

  Assumption 1 *Per-period payoffs* $\pi^{i,j}(x_t, \omega_t)$ *can be written as*

$$\begin{cases} \pi^{i,i}(x_t^i, \omega_t), & j = i, \\ \pi^{i,j}(\omega_t), & j \neq i. \end{cases}$$

- Law of motion: State-to-state transitions are controlled entirely by the player with the move.

  Assumption 2 *The law of motion is*

$$\Pr\left(\omega' | \omega_t, x_t\right) = \begin{cases} \Pr^i\left(\left(\omega'\right)^i | \omega_t^i, x_t^i\right), & \left(\omega'\right)^{-i} = \omega_t^{-i}, \\ 0, & \text{otherwise.} \end{cases}$$

Alternating-Moves Model

- Bellman equation if player i has the move ($i = j$):

$$V^{i,i}(\omega) = \max_{x^i} \pi^{i,i}(x^i, \omega) + \sqrt[N]{\beta} \mathsf{E}_{k',(\omega')^i} \left\{ V^{i,k'} \left(\left(\omega'\right)^i, \omega^{-i} \right) | \omega^i, x^i \right\},$$

  where k' denotes the next player to move.

- The strategy of player i is

$$X^i(\omega) = \arg\max_{x^i} \pi^{i,i}(x^i, \omega) + \sqrt[N]{\beta} \mathsf{E}_{k',(\omega')^i} \left\{ V^{i,k'} \left(\left(\omega'\right)^i, \omega^{-i} \right) | \omega^i, x^i \right\}.$$

- Bellman equation if player i does not have the move ($j \neq i$):

$$V^{i,j}(\omega) = \pi^{i,j}(\omega) + \sqrt[N]{\beta} \mathsf{E}_{k',(\omega')^j} \left\{ V^{i,k'} \left(\left(\omega'\right)^j, \omega^{-j} \right) | \omega^j, X^j(\omega) \right\}.$$

- If each of N players can move to one of K states, then expectation over successor states involves KN terms.

Alternating-Moves Model

- Problem: There are N Bellman equations per player.
- Solution: Define the expected value function

$$\overline{V}^i(\omega) = \frac{1}{N} \sum_{j=1}^{N} V^{i,j}(\omega).$$

$\overline{V}^i(\omega)$ is the value function *before* it is known who is next to move.

- The Bellman equation of player i is

$$\begin{aligned} \overline{V}^i(\omega) \quad = \quad & \frac{1}{N} \Bigg\{ \max_{x^i} \pi^{i,i}(x^i, \omega) + \sqrt[N]{\beta} \mathsf{E}_{(\omega')^i} \left\{ \overline{V}^i \left(\left(\omega'\right)^i, \omega^{-i} \right) | \omega^i, x^i \right\} \\ & + \sum_{j \neq i} \pi^{i,j}(\omega) + \sqrt[N]{\beta} \mathsf{E}_{(\omega')^j} \left\{ \overline{V}^i \left(\left(\omega'\right)^j, \omega^{-j} \right) | \omega^j, X^j(\omega) \right\} \Bigg\}. \end{aligned}$$

- The strategy of player i is

$$X^i(\omega) = \arg\max_{x^i} \pi^{i,i}(x^i, \omega) + \sqrt[N]{\beta} \mathsf{E}_{(\omega')^i} \left\{ \overline{V}^i \left(\left(\omega'\right)^i, \omega^{-i} \right) | \omega^i, x^i \right\}.$$

- Remaining problem: Discount factor is $\sqrt[N]{\beta}$ instead of $\beta \rightarrow$ iteration penalty.

Time to Convergence

#firms	simultaneous moves (mins.)	alternating moves (mins.)	time per iteration	ratio number of iterations	time to convergence
2	1.41(-4)	1.00(-4)	2.55	0.55	1.40
3	1.33(-3)	8.11(-4)	4.58	0.36	1.64
4	1.24(-2)	4.18(-3)	10.72	0.28	2.96
5	1.02(-1)	1.99(-2)	22.61	0.23	5.13
6	7.74(-1)	6.89(-2)	57.84	0.19	11.23
7	5.19(0)	2.03(-1)	150.00	0.17	25.54
8	3.21(1)	5.99(-1)	352.17	0.15	53.66
9	1.94(2)	1.48(0)	960.78	0.14	131.12
10	*1.10(3)*	3.43(0)	2,642.86	*0.12*	*320.59*
11	*5.80(3)*	7.37(0)	7,134.15	*0.11*	*786.81*
12	*2.94(4)*	1.53(1)	19,008.97	*0.10*	*1,923.53*
13	*1.51(5)*	2.97(1)	54,527.14	*0.09*	*5,101.20*
14	*7.08(5)*	5.62(1)	144,763.30	*0.09*	*12,592.71*

Time to convergence. (k) is shorthand for $\times 10^k$. Convergence criterion is "distance to truth$< 10^{-4}$." Entries in italics are based on an estimated 119 iterations to convergence in discrete time. Quality ladder model with $M = 9$ quality levels per firm and a discount factor of 0.925.

Equilibrium and Dynamics

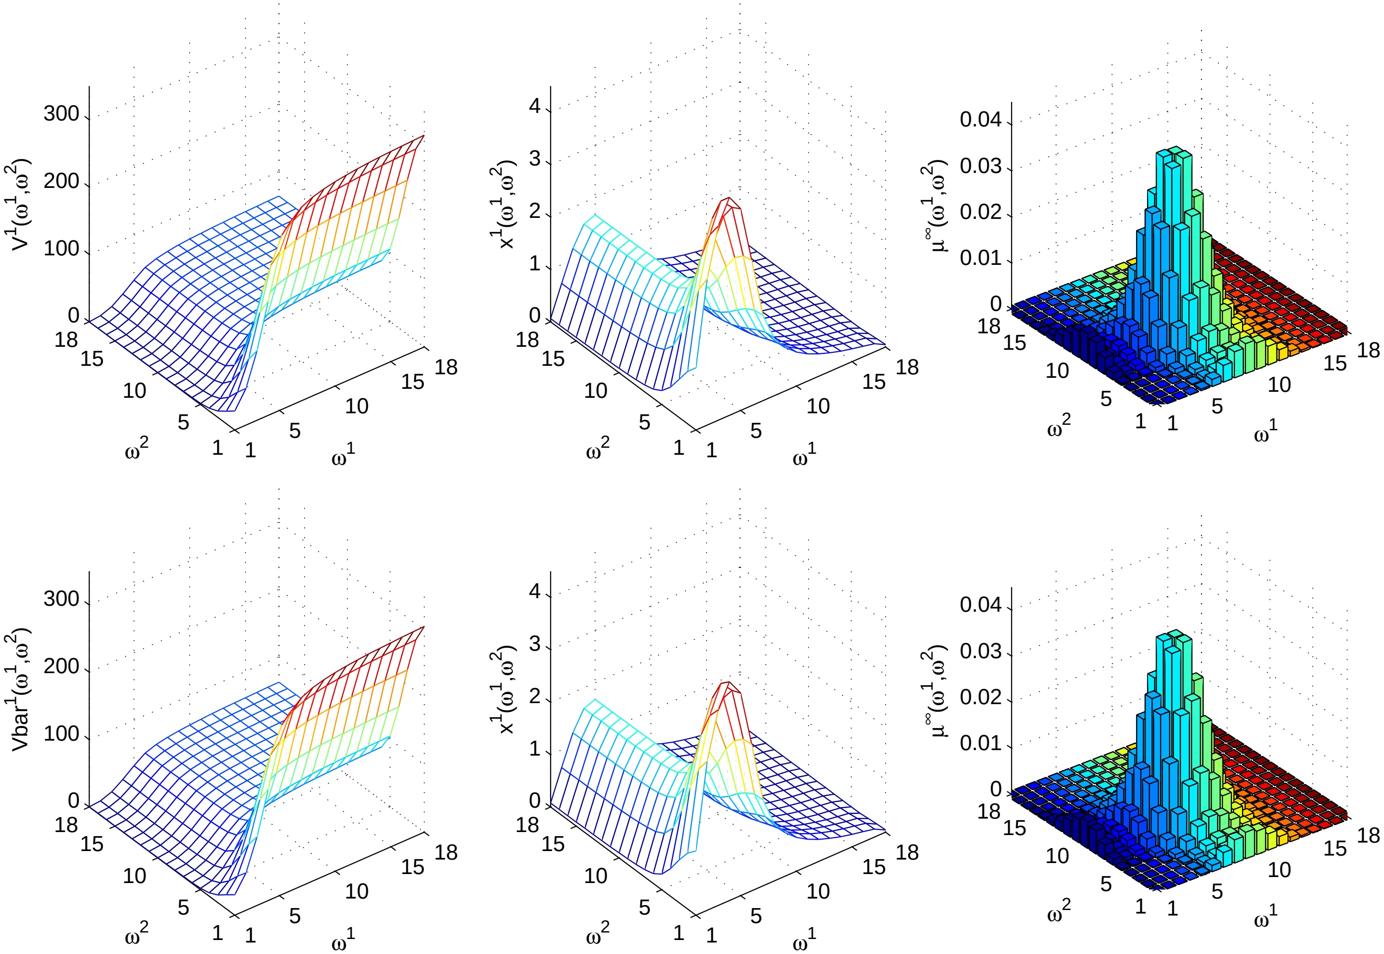

Equilibrium value and policy functions (left and middle panels) and limiting distribution (right panels). Simultaneous (upper panels) and alternating moves (lower panels). Quality ladder model with $N = 2$ firms, $M = 18$ quality levels per firm, and a discount factor of $\beta = 0.925$.

Open Questions

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

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."